Decoding Sustainable Investment Strategies: Bridging Intentions and Outcomes

Keer Yang[a] Ayako Yasuda[b]

April 2026

Abstract

We study sustainability intent of U.S. mutual funds and whether stated objectives translate into differences in behavior. Using machine learning applied to fund prospectuses, we classify sustainable strategies as financial, moral, or impact-focused. Among 1,523 funds managing $1.7 trillion in 2023, 88% of assets are financially motivated, 10% morally motivated, and just 2% impact oriented. Financial funds tilt toward firms with high ESG scores; moral funds rely on exclusionary screens and exhibit low flow–performance sensitivity. Only impact funds are associated with reductions in portfolio firms' carbon intensity. Most U.S. sustainable capital is therefore not deployed toward real-economy impact.

Keywords: Impact investing; Value vs. values; Machine learning; Carbon emissions; Proxy voting; Climate finance; ESG; SRI; Ethical divestment; Financial materiality; Sustainability

---

We are grateful for the comments of Vikas Agarwal, Aleksandar Andonov, Brad Barber, Ioannis Branikas, Hoa Briscoe-Tran, Murillo Campello, Will Cong, Yaniv Grinstein, Deeksha Gupta, Galina Hale, Po-Hsuan Hsu, Xing Huang, Wei Jiang, Laura Lindsey, Pedro Matos, David Ng, Nadya Malenko, Marcus Opp, Kelly Posenau, Aneesh Raghunandan, Yuan Ren, Alberto Rossi, Mirco Rubin, Tao Shu, Kelly Shue, Paul Smeets, Denis Sosyura, Laura Starks, Artem Streltsov, Sunil Wahal, Xiang Yin, Adam Zhang, and conference and seminar participants at AFA (Philadelphia), GSU CEAR-Finance Conference, CAFIN Conference, GRASFI Conference, Minnesota Family and Friends Conference, NTHU-UNSW-SMU Symposium, SGFIN Annual Research Conference on Sustainability, FMA, Arizona State University, Audencia Business School, Cornell University, EDHEC, Tsinghua University, UC Davis, University of Amsterdam, University of Tokyo, University of Essex, York University, and Zhejiang University. Sayar Banerjee and Yan Jia provided valuable research assistance. We each acknowledge the generous funding from the UC Davis Academic Senate Committee on Research and the Center for Investor Welfare. All errors are our own.

[a] University of California, Davis, Graduate School of Management; E-mail: kkeyang@ucdavis.edu
[b] University of California, Davis, Graduate School of Management; E-mail: asyasuda@ucdavis.edu

1. Introduction

Sustainable investing has grown rapidly, yet much of the literature and policy debate still abstracts from heterogeneity in the objectives pursued under the sustainable investing label. In practice, sustainable funds may seek to enhance financial performance through ESG integration, align portfolios with investors' moral or ethical values, or induce real-world change at portfolio companies.[1] Although recent theory and household-level evidence increasingly recognize heterogeneity in sustainable investment motives, much less is known about heterogeneity in the objectives of delegated intermediaries that market sustainable strategies and allocate capital on investors' behalf. We address this gap by developing a scalable framework that classifies mutual funds by sustainability objective—financial, moral, or impact—and use it to examine whether fund actions and outcomes align with those stated goals.

From a policy perspective, achieving environmental and social goals through private capital markets requires that funds be genuinely impact-oriented, that households and institutions have preferences aligned with those goals[2], and that capital be effectively matched to the appropriate funds. More generally, distinguishing among financial, moral, and impact objectives is essential for evaluating both the promise and the limitations of sustainable investing. Without such distinctions, it is difficult to interpret the growth of sustainable funds, assess whether fund behavior aligns with stated objectives, or determine whether sustainable capital is likely to generate real-world impact.

Despite the rapid growth in sustainable fund assets under management (AUM), relatively little is known about the objectives underlying these investment strategies or whether fund behavior is consistent with those objectives. Existing classification systems used by practitioners,

[1] Theoretical and empirical work on sustainable investing studies nonpecuniary demand for sustainable assets and, in some cases, distinguishes between financially and non-financially motivated sustainable investing; see, e.g., Fama and French (2007), Benabou and Tirole (2010), Pedersen, Fitzgibbons, and Pomorski (2021), Pástor, Stambaugh and Taylor (2021), Oehmke and Opp (2025), and Green and Roth (2025).

[2] A related strand of work examines investors' willingness to accept lower financial returns in exchange for investing in funds that pursue nonfinancial goals, while more recent work further distinguishes between impact-oriented and deontological or expressive motives; see, e.g., Riedl and Smeets (2017), Barber, Morse, and Yasuda (2021), Bauer, Ruof, and Smeets (2021), Heeb et al. (2023), and Bonnefon et al. (2025).

regulators, and researchers often fail to distinguish among fundamentally different sustainability goals.[3] In the U.S. in particular, where regulatory frameworks remain limited, commonly used sustainable investing labels often pool together funds pursuing financial, moral, and impact objectives. As a result, empirical studies and policy debates often rest on untested assumptions about what sustainable funds are trying to achieve.

This ambiguity has contributed to competing narratives about sustainable investing. One common critique is that sustainable funds pursue ideological goals at the expense of shareholder value, implicitly assuming that many such funds are designed to induce real-world change and therefore accept weaker financial performance. A contrasting view is that the limited evidence of real-world impact is disappointing, again assuming that many sustainable funds are intended to reduce negative externalities. Both critiques rest on a shared but largely unexamined premise: that a substantial share of sustainable funds are impact-oriented. Whether that premise holds in practice is an open empirical question.

Recent research increasingly distinguishes between financially motivated and nonfinancially motivated sustainable investing, underscoring the importance of separating sustainability objectives grounded in financial materiality from those reflecting investors' nonpecuniary preferences (e.g., Starks (2023)). However, much of this progress has focused on investor motives rather than on the delegated intermediaries that market sustainable strategies and allocate capital on investors' behalf. An important empirical challenge therefore remains: how to identify, at scale, the objectives that sustainable funds themselves pursue. We address this challenge by developing a scalable text-based framework that classifies mutual funds according to their stated sustainability objectives—financial, moral, or impact.

We begin by distinguishing three core motivations for sustainable investing: financial, moral, and impact. Financially motivated funds seek to improve portfolio performance by integrating ESG-related risks and opportunities, treating sustainability characteristics as inputs to

[3] Related work highlights ambiguity in ESG measurement and labeling, including disagreement in firm-level ESG ratings and incomplete alignment between broad ESG labels and underlying fund behavior; see, e.g., Berg, et al. (2022), Christensen et al. (2022), Raghunandan and Rajgopal (2022), and Allcott et al. (forthcoming).

stock selection or risk management.[4] Morally motivated funds aim to align portfolios with investors' ethical values, often through exclusions of sectors such as tobacco, fossil fuels, or weapons. Impact-oriented funds seek to generate measurable improvements in environmental or social outcomes, including through ownership and engagement.

This distinction between moral and impact investing is especially important. Both reflect nonpecuniary preferences, but they operate through different mechanisms. Moral investors seek to avoid complicity in harmful activities, whereas impact investors seek to improve firm behavior and real-world outcomes through capital allocation and stewardship. For example, a moral fund may exclude a high-emissions utility, whereas an impact fund may overweight the same firm with the aim of reducing its emissions over time through engagement.[5] These differences imply distinct portfolio compositions, forms of stewardship, and expected outcomes.

In this paper, we investigate the prevalence, behavior, and effectiveness of three types of sustainable mutual funds—financial, moral, and impact. We ask three related questions: What proportions of U.S. mutual funds pursue these different sustainable investing strategies? Do funds act consistently with their stated sustainability goals? And which types of fund investments, if any, are associated with real-world environmental or social improvements? To answer these questions, we develop a fine-tuned BERT model that identifies whether a fund frames sustainability in financial, moral, or real-world impact terms. This yields a fund-level measure of sustainability intent, which we validate on a labeled sample and then use to examine differences in portfolio composition, fund performance, investor flows, stewardship behavior, and subsequent firm-level emissions.

Applying this framework, we identify more than 1,500 sustainable mutual funds managing \$1.7 trillion in assets as of 2023, up fifteen-fold from \$130 billion in 2014, when just

---

[4] Our distinction between financially and nonfinancially motivated sustainable investing reflects a broader conceptual distinction in the literature between sustainability considerations grounded in financial materiality and those grounded in nonpecuniary preferences; see, e.g., Pedersen, Fitzgibbons, and Pomorski (2021) and Starks (2023).

[5] Our distinction between moral and impact motives is related to recent work distinguishing among different nonpecuniary objectives in sustainable investing; see, e.g., Oehmke and Opp (2025), Heeb et al. (2023), and Bonnefon et al. (2025).

under 150 such funds existed. Yet even by 2023, sustainable AUM remains overwhelmingly concentrated in financially motivated funds: 88% of assets are managed by financial funds, compared with 10% by moral funds and 2% by impact funds. In dollar terms, impact funds account for only $29 billion, or 0.08% of the total U.S. mutual fund market. We further show that much of the recent growth in sustainable AUM—especially among financial funds—reflects newly classified funds rather than only the expansion of previously classified funds. Supporting this interpretation, changes in fund classification are accompanied by meaningful changes in prospectus language, suggesting that reclassification reflects shifts in stated sustainability objectives rather than merely generic ESG wording or broader time trends.

The observed distribution of fund types likely reflects the joint influence of investor demand, fund-manager incentives, and frictions in the search and matching process. Households and institutions may value different sustainability goals[6], fund managers may respond to that demand subject to their own incentives and constraints, and imperfect transparency or investor misunderstanding may further affect how capital is allocated across fund types. While we do not separately identify these drivers, the skewed distribution of capital is itself informative: absent regulatory or structural changes, most capital labeled as sustainable is unlikely to be deployed toward generating real-world impact.

We show that sustainability intent predicts systematic differences in portfolio composition and fund characteristics. Financial funds tend to be larger and hold firms with high ESG ratings and low carbon intensity, consistent with ESG integration for risk management. Moral funds systematically exclude sin sectors and have lower expense ratios. Impact funds are smaller, have higher expense ratios, and overweight high-emissions sectors, consistent with an engagement-oriented strategy.

We next examine whether sustainability intent is associated with differences in fund performance, investor behavior, and stewardship. While we find no meaningful differences in risk-adjusted performance across fund types, sustainability intent predicts systematic differences

[6] See, e.g., Krueger, Sautner, and Starks (2020), Barber, Morse, and Yasuda (2021), and Giglio et al. (2025).

in investor flows and voting behavior. Moral funds exhibit lower flow–performance sensitivity, consistent with a less performance-driven investor base. Impact funds are more likely to support outcome-oriented environmental and social shareholder proposals, both relative to financial and moral funds and relative to other types of environmental and social proposals.

Finally, to assess whether sustainability intent is associated with real-economy outcomes, we use a difference-in-differences design to examine whether firms newly added to the portfolios of different fund types exhibit subsequent changes in carbon intensity relative to matched firms. We find that firms added by impact funds subsequently reduce carbon intensity, whereas no comparable pattern emerges for firms added by financial or moral funds. Yet because impact funds manage only a very small share of assets, their aggregate influence remains limited. More broadly, our findings suggest that once sustainable funds are distinguished by objective, their behavior is largely consistent with those objectives. The central limitation is not that all sustainable funds fail to do what they claim, but that most capital labeled as sustainable is not allocated to funds seeking measurable real-economy change.

Contribution and Related Literature. This paper makes several contributions to the literature on sustainable investing and mutual funds. First, we contribute to the literature on sustainable investing by classifying mutual funds according to sustainability intent—financial, moral, or impact—using the fund's own prospectus language. Prior research shows that investors in sustainable investment products differ in their motivations, willingness to accept lower returns, and demand for sustainability characteristics. In mutual funds and related public-market settings, investors exhibit substantial heterogeneity in motives for sustainable investing (Riedl and Smeets (2017), Baker, Egan, and Sarkar (2022), Giglio et al. (2025)). Related evidence from delegated pension settings and private-market impact investing further suggests that some investors are willing to accept lower financial returns in exchange for nonpecuniary objectives (Bauer, Ruof, and Smeets (2021), Barber, Morse, and Yasuda (2021)). We build on this literature by showing that the broad category of sustainable mutual funds itself contains economically meaningful heterogeneity in stated objectives. In related work, Lowry, Wang, and Wei (2025) distinguish among ESG funds based on incentives to engage, capturing variation in financially motivated sustainability. Our approach differs in using fund text to identify not only financially motivated sustainability intent, but also to distinguish moral and impact objectives within the broader set of nonfinancially motivated funds.

Second, the paper contributes to the growing literature that uses machine learning and textual analysis to measure sustainability-related constructs in finance.[7] Prior work documents substantial disagreement across ESG rating providers and the challenges this creates for research and practice (Berg et al. (2022), Christensen et al. (2022)). These concerns, together with recent advances in natural language processing, have motivated researchers to develop new sustainability measures based on textual data.[8] In mutual funds research, Abis (2022) and Abis and Lines (2024) use unsupervised topic models to classify funds' overall investment strategies. We contribute to this literature by applying a fine-tuned BERT model to mutual fund prospectuses to identify sustainability intent. In this respect, our approach is also related to Rajan, Ramella, and Zingales (2022), who use NLP to recover stated corporate objectives from shareholder letters. We compare a fine-tuned BERT model with generative-AI approaches using a manually labeled sample and find that the BERT-based model performs better for our application.

Applying this classification, we find that the majority of funds using sustainability language articulate primarily financial objectives, with relatively few express a distinct impact-oriented mandate. These findings complement contemporaneous work by Abis, Buffa, and Sadashivam (2024), who manually classify mutual funds' sustainability strategies based on prospectuses language, and by Edmans, Gosling, and Jenter (2026), who survey institutional investors about their ESG objectives. Relative to these studies, we construct a manually labeled sample, use it to assess the BERT model's performance, and then apply the model at scale to the universe of U.S. mutual funds. We then link these classifications to fund behavior, stewardship, and firm-level outcomes.

---

[7] See Gentzkow, Kelly, and Taddy (2019) for a survey of textual analysis in economics and finance, and Jiang, Wang, and Yang (2025) for a review of AI-based ESG research.

[8] See, e.g., Engle et al. (2020), Sautner et al. (2023), Duchin et al. (2025), and Briscoe-Tran (2026).

Third, we contribute to the literature on investor demand, fund behavior, and stewardship in sustainable investing. Prior work shows that sustainable mutual funds exhibit lower flow-performance sensitivity than conventional funds and that this relation can vary across sustainability screens and fund characteristics (e.g., Bollen (2007), Renneboog, ter Horst, and Zhang (2011), and Bialkowski and Starks (2018)). Hartzmark and Sussman (2019) document that the introduction of Morningstar sustainability ratings led investors to reallocate toward high-sustainability funds and away from low-sustainability funds, while Gantchev, Giannetti and Li (2024) show that this response is at least partly driven by expectations of superior financial performance rather than purely nonpecuniary motives. We add to this literature by showing that lower flow-performance sensitivity is most clearly concentrated in funds with moral sustainability intent, a pattern consistent with nonpecuniary investors sorting into moral funds. For funds with financial or impact objectives, the corresponding estimates are imprecise, suggesting weaker or noisier differentiation in investor clientele.[9]

Another strand of literature examines engagement and shareholder voting on environmental and social issues.[10] Much of this work focuses either on investors whose objectives are primarily financial or on sustainable investors without explicitly distinguishing among financial, moral, and impact motives. Relative to this literature, our contribution is to compare voting behavior across funds with distinct stated sustainability objectives. We show that all three types of sustainable funds are more likely than conventional funds to support environmental and social shareholder proposals, but that support is strongest among impact funds, especially for proposals aimed at changing firm outcomes than expanding disclosure.

Fourth, we contribute to the literature linking sustainability objectives to real-economy outcomes. Lowry, Wang, and Wei (2025) study heterogeneity within financially motivated ESG

---

[9] This interpretation is consistent with recent evidence that retail investors' preferences may be more closely aligned with moral funds than with impact funds (Bonnefon, Sastry, and Thesmar (2025); Heeb et al. (2023)). Relatedly, Heeb, Kölbel, and Weder (2025) suggest that retail investors may overestimate the impact of generic sustainable products.

[10] See, e.g., Dimson, Karakas, and Li (2015, 2025), Flammer et al. (2021), Heath et al. (2023), He, Kahraman, and Lowry (2023), and Couvert (2025).

funds, showing that funds with stronger incentives to engage have greater effects on portfolio firms' ESG risk and emissions following severe ESG incidents. Our contribution differs in shifting the focus from heterogeneity within financially motivated ESG funds to heterogeneity across financial, moral, and impact objectives. Our framework is also related to Oehmke and Opp (2025), whose model distinguishes between investors who seek to reduce negative externalities through engagement-oriented investments and those who prefer to hold already greener firms.[11] Our empirical results align with this distinction in several ways. First, we show that impact funds systematically hold larger portfolio weights in high-emission sectors such as utilities. Second, we show that impact funds are more likely than financial and moral funds to support environmental and social shareholder proposals, particularly those aimed at changing firm outcomes. Finally, we show that firms held by impact-oriented funds subsequently reduce emission intensity relative to matched firms, whereas no such pattern is observed for financial or moral funds. These findings suggest that real-economy change is concentrated in funds whose stated objective is to influence firm behavior, rather than in funds whose primary goal is to hold greener portfolios.

The remainder of the paper proceeds as follows. Section 2 describes the data and our classification of fund sustainability intent. Section 3 presents the distribution and evolution of fund sustainability intent and examines how it maps into portfolio composition and ESG-related firm characteristics. Section 4 analyzes differences in fund performance, investor flows, and voting behavior. Section 5 studies whether sustainability intent is associated with firm-level emissions outcomes. Section 6 concludes.

[11] Related work also shows that the real-economy effects of exclusionary or portfolio-tilting strategies depend on the underlying mechanism and market environment; see, e.g., Hartzmark and Shue (2023), Dasgupta et al. (2025), Gupta et al. (2026), and Dangl et al. (forthcoming). In related empirical work, Atta-Darkua et al. (2023) find that institutional investors joining climate pledges reduce their portfolio emissions primarily through rebalancing toward low-emission stocks rather than by engagement.

2. Empirical Method and Data

This section describes our empirical methodology and data. Section 2.1 outlines our classification of sustainability intent. Section 2.2 details the text-based methodology used to implement this classification across U.S. actively managed mutual funds and ETFs. Section 2.3 describes the additional data sources used in the analysis.

2.1 Conceptual Framework for Sustainability Intent

Building on the framework introduced in Section 1, we classify sustainability intent along two dimensions. First, we distinguish between value-driven (financial) investing, in which ESG characteristics are incorporated insofar as they affect expected returns or risk, and values-driven investing, in which investors derive utility directly from environmental or social characteristics. Second, among values-driven investors, we distinguish between categorical morality and impact objectives. Together, these dimensions yield three sustainability goals: financial value, categorical morality, and impact.

Financial funds integrate ESG information to improve risk-adjusted performance. In this case, ESG characteristics matter instrumentally—as inputs into financial decision-making—rather than intrinsically.

Moral funds pursue categorical exclusion based on ethical principles, such as avoiding specific industries or activities. Impact funds seek to generate measurable improvements in environmental or social outcomes through investment and engagement. Although both reflect nonpecuniary preferences, their mechanisms and implied portfolio strategies differ.

This framework provides a disciplined basis for classifying sustainable funds by stated intent and for generating testable predictions about portfolio composition, investor behavior, stewardship, and real-economy outcomes. In Section 2.2, we operationalize this framework using a text-based classification of mutual fund prospectus disclosures.

2.2 Text-Based Classification of Sustainability Intent

There are no formal U.S. regulations that define sustainable funds or distinguish among financial, moral, and impact objectives. As a result, investors must rely on their own research or on third-party providers to identify funds aligned with their sustainability goals. We therefore develop a text-based empirical method to classify funds by stated sustainability intent.

2.2.1 Prospectus-Based Classification Procedure

We extract the "Principal Investment Strategy" section of the prospectus, as this required disclosure describes how the fund "intends to achieve its investment objective" (U.S. SEC, 2016) and typically contains statements regarding sustainability intent. Prospectuses (Form 497K and/or Form 485) are obtained from the SEC EDGAR system for 2014 to 2023. Because this section also describes non-sustainability objectives, we first use an ESG keywords list to isolate ESG-related sentences that discuss ESG-related topics and discard the rest.[12] This yields 75,886 fund-quarter observations.[13]

We train a supervised machine-learning model using a manually labeled sample of prospectus sentences and apply it to classify the full text corpus consistently and at scale. Because the three sustainability goals differ in how investors value non-financial traits and outcomes—not merely in the topics discussed—it is critical to capture investor intent rather than rely solely on keywords. We therefore employ a sentence-level method that interprets contextual meaning rather than isolated word pairs or topic nouns.

For this purpose, we use a Bidirectional Encoder Representations from Transformers (BERT) model. BERT is a natural language processing (NLP) architecture released by Google in 2018 whose key feature is the ability to interpret words in relation to surrounding context within a sentence or passage. The model is pre-trained on large corpora of text provided by Google and can be fine-tuned by researchers with smaller labeled datasets for specific NLP tasks such as classification.

In our setting, we fine-tune BERT to classify sentences from fund prospectuses as "financial value," "moral (categorical) investing," or "impact investing". We then aggregate these sentence-level classifications to construct fund-level measures of sustainability intent.

---

[12] Our ESG keyword list is intentionally broad; sentences in the Principal Investment Strategy section without ESG-related keywords are treated as unrelated to sustainability. The complete list is provided in Appendix A1.
[13] Each fund–quarter is matched to its most recent Principal Investment Strategy disclosure.

To train and evaluate the model, we construct a labeled sample of ESG-related sentences that clearly express each of the three sustainability goals. From the 75,886 fund-quarter observations, we sample 362 observations and identify 3,575 ESG-related sentences for manual classification. Because the underlying data are highly imbalanced, we oversample true positives (sentences expressing financial, moral, or impact goals) to improve classification balance and performance, following common practice in machine learning (He and Garcia, 2009).

Of the 3,575 sentences used for training and testing, 2,438 are drawn from the prospectuses of a random subset of U.S. sustainable mutual funds compiled by Morningstar. We oversample this group because these funds are more likely to explicitly describe their sustainability goals in the investment strategy section of the prospectuses. We supplement the training and testing sample with sentences drawn from funds not included in the Morningstar list.[14] All sentences are manually classified by the authors as “financial”, “moral”, “impact”, or “unclassified”.

We classify a sentence as “financial value” if it indicates that ESG (or other non-financial) information is used to improve financial performance. For example: “The Adviser evaluates financially material environmental, social, and governance (ESG) factors as part of the investment decision-making process for the Fund.” This sentence is classified as financial value because ESG factors are evaluated for their relevance to financial performance.

We classify a sentence as “moral (categorical) investing” if it indicates that certain investments are excluded based on ethical considerations. For example: “The Index excludes companies involved with weapons, tobacco, gambling, alcohol, adult entertainment, and nuclear power.” This sentence is classified as moral because the exclusions target traditionally defined “sin” industries. In some cases, funds also cite religious or faith-based principles as the basis for exclusion.

[14] There is no fixed rule for the appropriate size of a training and testing sample in this setting. When selecting fund-quarter observations for manual labeling, we include each fund only once to maximize cross-sectional variation in sentence structures and phrasing rather than repeatedly sampling similar language from the same fund over time. More details on the construction of the training and testing sample are provided in Internet Appendix Section IA2.

We classify a sentence as “impact investing” if it indicates that ESG information is used to pursue positive externalities. For example: “Under normal market conditions, the Fund will invest at least 80% of its net assets … in the equity securities of companies identified by the Fund’s investment sub-advisor, … as having or seeking to have a positive carbon impact.” This sentence is classified as impact because investments are selected for their intended environmental benefit.

We label 729 sentences as “financial”, 288 sentences as “moral”, and 305 sentences as “impact”.[15] Among the 3,575 manually classified sentences, we randomly split the data into an 80% training sample and a 20% testing sample using stratified sampling to preserve balanced class representation.

The BERT model is trained on the training sample to learn sentence patterns corresponding to financial value, moral (categorical) investing, and impact investing. The three categories are trained separately. We then use the model to classify sentences in the testing sample as “financial”, “moral”, and “impact”. Because the true labels are known, we evaluate model performance by comparing the model’s predictions with the manual classifications using three standard metrics: accuracy, precision, and recall. Table 1 reports the results.

We find that the BERT model performs well in classifying the sustainability goal expressed in a sentence. Accuracy ranges from 91% for financial sentences to 98% for moral sentences, meaning that the model correctly classifies most sentences as either positive or negative cases. Precision ranges from 83% for impact to 85% for moral sentences, meaning that most sentences classified as positive cases are true positives. Recall ranges from 66% for impact to 89% for moral sentences, meaning that the model correctly identifies most of the true positive cases. The f1 score ranges from 73% for impact to 87% for moral sentences and summarizes the balance between precision, which penalizes false positives, and recall, which penalizes false

[15] Most sentences (2,278) are classified as “unclassified,” either because they contained ESG-related keywords (e.g., “transform”) without expressing a sustainability objective or because the stated objective is too vague or ambiguous to categorize. In addition, 25 sentences (0.7%) meet the criteria for two of the three categories (e.g., both moral and impact). Additional details on the classification criteria, including example sentences for each of the four manually coded categories, are provided in Internet Appendix Section IA1.

negatives. These performance levels are comparable to those reported in other studies using BERT models (e.g., Bingler et al. (2024) and Rajan et al. (2022)).

After establishing the reliability of our BERT model, we apply it to all sentences from the 75,886 fund–quarter observations and construct fund–quarter–level measures based on the sentence-level classifications. Each fund is assigned the category with the largest percentage share of sustainable sentences, which we interpret as the fund's primary sustainable goal. Funds may also have secondary motives, reflected in smaller numbers of sustainable sentences from other categories.

When there is a two-way (or three-way) tie for the largest percentage share, the fund is assigned to each of the tied categories. In such cases, the fund is interpreted as pursuing multiple sustainability goals as primary objectives. This occurs in 1,067 fund–quarter observations, representing less than 2% of the total sample.

2.2.2 Comparison to GenAI Models

An alternative classification approach is to use generative AI (GenAI) models such as ChatGPT, Google's Gemini series, or Claude. In principle, these models can be prompted to classify sustainable investment goals from prospectus text.

Two approaches are possible. The first is to prompt the model to perform the classification directly without providing training data. In this case, the model's pretrained knowledge implicitly serves as the ground truth rather than researcher-defined labels. This approach is problematic because prior research shows that retail investors and the media discussions often misunderstand sustainable investment goals (Hartzmark and Sussman (2019), Gantchev, Giannetti, and Li (2024), Heeb, Kölbel, and Weder (2025)). Such misunderstandings may also appear in the text corpora used to pretrain large language models and may therefore influence their classifications (Bini et al. (2026)). Internet Appendix Table IA2 shows that the ChatGPT-5 model performs worse than the BERT model in our classification task.

A second approach is to supply the training sample through prompting or retrieval-augmented generation (RAG) so that the model's outputs better align with researcher-defined classifications. However, because GenAI models generate text probabilistically, their outputs can vary across runs and may be difficult to reproduce exactly. When researcher-defined ground truth is available, a supervised BERT model offers greater stability and full reproducibility while avoiding hallucinations.[16] Internet Appendix Section IA3 provides additional discussion.

2.3 Additional Data Sources

Our empirical analysis draws on several additional data sources. We construct fund-quarter-level ESG portfolio characteristics by merging the CRSP Mutual Fund database with the MSCI ESG Ratings database. From CRSP, we obtain the average market value of each stock held by the fund for the quarter, which we use to compute portfolio weights. From MSCI, we obtain ESG ratings for the stocks held by the funds in our sample. Combining the two databases, we calculate fund-quarter-level ESG ratings as the value-weighted average ESG rating of all MSCI-rated stocks held by the fund. We use both the industry-adjusted and unadjusted ratings to address different hypotheses.

We construct fund-year-month-level fund flows and fund alphas from the CRSP Mutual Fund dataset. At the fund-year-month level, a fund's sustainability type is determined using the most recently available quarterly classification. We construct firm-month-level and fund-month-level Scope 1&2 and Scope 1&2&3 carbon emission intensity measures using the Trucost Environmental dataset, defined as metric tons of emissions per million dollars of revenue. Trucost reports firms' emission intensity annually. We convert these annual measures into firm-month emission intensity using a weighted average of the two adjacent fiscal years, with weights equal to 12 minus the distance (in months) from each fiscal year. We use both industry-adjusted and unadjusted emission intensity measures. Fund-level emission intensity is constructed as the value-weighted average of firm-level emissions across all portfolio stocks with non-missing emission data.

Finally, we obtain funds' voting records on shareholder proposals from the ISS database. We match ISS database FundIDs to the FundIDs in our sample using the methodology of

[16] Huang et al. (2023) discuss a similar limitation of generative language models in financial text applications and document improved stability and reproducibility when supervised models are trained on researcher-labeled data.

Sulaeman and Ye (2023). Specifically, we use SEC Form N-PX and the reference identifier (NPXFileID) provided by ISS to link each fund to its corresponding filing. We then use name matching to identify the fund identifier (series_id) in Form N-PX and merge it with our dataset. We construct a fund-level ESG activism measure defined as the likelihood of supporting shareholder proposals on environmental and social issues.

3. Sustainable Fund Universe and Portfolio Composition

3.1 Distribution and Growth of the U.S. Sustainable Funds Universe

Using our empirical method, we classify U.S. sustainable funds into financial, moral, and impact categories. Figure 1 presents the fund count, AUM levels and the sources of asset growth of sustainable funds by fund type over the sample period.

We find that 1,523 sustainable funds managed $1.7 trillion in 2023, compared to 145 sustainable funds managing $0.1 trillion in 2014 (Panel A-1 and A-2). Relative to the CRSP mutual fund universe, sustainable funds' share of AUM increases from 0.5% to 4.7% (shown in the Internet Appendix Figure IA1). Overall, sustainable funds are fast-growing but remain a minor segment of the mutual fund industry in the U.S.

Among the three categories, financial funds dominate, and their share of sustainable fund AUM increases over time. In 2014, financial funds managed 82% of total sustainable AUM (49% by fund count, or 71 of 145 funds,). By 2023, they managed 88% of total AUM (80% by fund count, or 1,211 of 1,523 funds). The AUM managed by financial funds grew 14 times from $108 billion to $1.470 trillion, while impact funds' AUM grew only 5 times from $6 billion to $29 billion during the sample period. Even under a more lenient definition, impact-oriented fund AUM remains small.[17] Moral funds fall in between, with AUM increasing tenfold from $18

---

[17] In 2023, under the most lenient classification—requiring only one impact-related sentence in the prospectus—the combined AUM of such funds is $86 billion. As impact is not necessarily the dominant objective for these funds, this amount (about 5% of total sustainable fund AUM) represents an upper bound on impact-oriented assets.

billion to $176 billion. To our knowledge, this overwhelming prevalence of financial value-driven sustainable funds has not been systematically documented in the literature.[18]

Although impact funds have grown in absolute terms, their influence has declined relative to financially driven sustainable funds. As a result, treating sustainable funds as a homogeneous category—common in the literature—risks obscuring substantial and time-varying heterogeneity. This heterogeneity operates along two margins: between *value* and *values* objectives, and, within values-driven funds, between *moral* and *impact* strategies.

Our methodology identifies a broader and more systematically defined set of sustainable funds than third-party providers such as Morningstar, which rely on proprietary criteria.[19] As of year-end 2021, Morningstar reports $357 billion in sustainable fund AUM, compared to $1.5 trillion under our classification.

The difference reflects both scope and transparency. For example, two Morgan Stanley funds with similar prospectus language regarding financially motivated ESG engagement are treated differently by Morningstar. The Morgan Stanley Growth Portfolio ($16.1B TNA), which describes ESG engagement as "material to the value of the security over the long term"[20] is included in our sample but not in Morningstar's, whereas the smaller Morgan Stanley Global Sustain Portfolio ($113M TNA), which refers to ESG factors as "material to long-term sustainably high returns"[21] appears in both.

Because vendor methodologies are proprietary and subject to revision, researchers lack visibility into the criteria underlying such distinctions. In contrast, our text-based classification applies uniform, observable rules across funds and over time, enabling transparent and replicable empirical analysis.

---

[18] See also Abis, Buffa, ad Sadasivam (2024) for contemporaneous and complementary evidence on variation in ESG fund strategies.

[19] Third-party sustainable fund lists rely on proprietary criteria that may change over time.

[20] https://www.sec.gov/Archives/edgar/data/836487/000113322821002721/msif-html3680_497k.htm

[21] https://www.sec.gov/Archives/edgar/data/836487/000113322821002712/msif-html3679_497k.htm

Our approach also yields a more precise benchmark than PRI signatory AUM, which mechanically aggregates the full AUM of signatory firms. PRI figures therefore overstate sustainable assets by including capital allocated to non-sustainable strategies. For instance, BlackRock—whose $5 trillion in equity AUM ($2.2 trillion in CRSP equity funds) as of 2023 reflects investor demand across diverse mandates—is a PRI signatory. Yet only $52 billion of its CRSP equity fund AUM is classified as sustainable under our method. This gap highlights the distinction between asset manager-level commitments and the actual investor allocation to sustainable strategies.

Panels B-1 to B-3 decompose total AUM within each category (financial, moral, and impact) into assets managed by funds classified in prior years ("Existing") and those managed by newly incepted or newly classified funds ("Entrants"). Across categories, existing funds account for the majority of AUM. Thus, the stock of sustainable assets primarily reflects the scale and expansion of incumbent funds rather than widespread entry. This pattern is especially pronounced for impact funds. Financial funds in 2021 and moral funds in 2022 are partial exceptions, when entrants account for a larger share of AUM; even in those years, however, existing funds remain central to the overall asset base.

Panels C-1 to C-3 shift the focus from levels to changes, decomposing year-over-year ΔAUM. Here, a different margin emerges. For financial funds, AUM growth during 2020–2023 is driven largely by funds that become newly classified as sustainable, even as existing funds experience net outflows in 2022. Moral funds also see meaningful contributions from newly classified funds, particularly in 2022. These dynamics are more muted for impact funds.

For financial funds during 2020–2023, increases in AUM are driven largely by funds that become newly classified as sustainable, even as existing funds experience outflows in 2022. Moral funds display a similar, though less pronounced, pattern—most notably in 2022—while these dynamics are comparatively muted for impact funds.

As shown in Figure 1, growth in sustainable fund AUM—particularly among financial funds—largely reflects previously non-sustainable funds that become newly classified as sustainable following changes in prospectus language. Figure 2 examines whether prospectus language changes meaningfully around these classification changes.

Panel A focuses on funds that switch classification status. Newly classified funds contain no sustainability sentences prior to classification by definition, but average 2.2 such sentences in the classification year. Existing sustainable funds exhibit little year-over-year change, averaging 2.4 sustainability sentences. Funds that are subsequently declassified display comparable levels prior to declassification, which then drop to zero.

ESG-keyword-only sentences—those containing ESG-related keywords but not classified as sustainability sentences—average 3.1 for existing sustainable funds and increase from 0.8 to 3.1 for newly classified funds. In contrast, Panel B shows that for non-sustainable funds whose classification does not change, ESG-keyword-only sentence counts remain stable over the same event window.

Taken together, these patterns indicate that prospectus language adjusts meaningfully when funds adopt or abandon sustainable strategies, and that the observed changes reflect shifts in stated sustainability objectives rather than general time trends in ESG disclosure.

Table 2 presents summary statistics for sustainable funds relative to benchmark funds. Beginning with Table 2, we restrict the sample to equity funds using CRSP asset class codes.[22] Benchmark funds are defined as CRSP equity non-index funds and ETFs whose prospectuses contain no sustainability keywords and that never appear on Morningstar's sustainable fund list during the sample period. The unit of observations is a fund-quarter from 2014 Q2 to 2023 Q4. Unless otherwise noted, subsequent analyses focus on equity funds.

Averaged across the 39 quarters, there are approximately 314 financial funds, 90 moral funds, and 46 impact funds, compared to 1,524 benchmark funds. All fund types allocate roughly 90% of assets to equity. Financial and impact funds hold more concentrated portfolios, with fewer than 100 stocks on average, whereas moral funds hold nearly 200 stocks—more than

[22] This restriction modestly reduces total sustainable fund AUM. In 2023, sustainable funds manage $1.7 trillion in total AUM, of which $1.2 trillion is invested in equity funds and $0.5 trillion in taxable and tax-exempt fixed income funds.

benchmark funds. Benchmark funds are larger on average than sustainable funds; among sustainable funds, financial funds are the largest ($1.226 billion) while impact funds are the smallest ($327M).

Expense ratios differ across categories. Impact funds charge the highest retail expense ratios (1.47%), consistent with more resource-intensive strategies. Moral funds, which primarily rely on ex ante screening, charge the lowest (1.27%). Financial funds' expense ratios (1.34%) are comparable to benchmark funds (1.35%).

Figure 3 compares sector weights across fund types and reports average sector-level E and S ratings and emission intensity (based on NAICS classifications). Financial funds' sector allocation closely track benchmark funds. In contrast, impact funds significantly overweight utilities and manufacturing and underweight information and finance. Utilities and manufacturing include firms with higher emission intensity and low E ratings, whereas information and finance have lower emission intensity and high E ratings.

This pattern is consistent with impact funds allocating capital to sectors with greater scope for environmental improvement rather than concentrating in already "green" industries (Hartzmark and Shue, 2023). We examine voting behavior and changes in firm-level environmental performance during holding periods in later sections.

We further examine exclusionary practices of moral funds. Table IA3 reports portfolio weights in sectors commonly excluded in prospectuses—tobacco, oil & gas extraction, natural gas distribution, coal mining, casino & gambling, and aerospace (weapons). Moral funds underweight these sectors relative to benchmark funds, consistent with stated exclusionary objectives.

3.2 ESG Characteristics and Sector Allocation

To assess whether portfolio holdings differ across fund types, we compare value-weighted ESG, E, and S ratings. Financial funds are expected to hold higher-rated firms, while impact funds may target firms with weaker ESG performance and greater scope for improvement.

Table 3 reports summary statistics and regression results. Columns (1) - (4) of Panel A report mean ESG ratings by fund type, showing that financial funds hold stocks with higher ESG, E, and S ratings than benchmark funds. Impact funds hold stocks with lower E ratings than financial funds. Because ESG scores trend upward over time and financial funds become more prevalent in later years, columns (5) and (6) include year–quarter fixed effects to isolate cross-sectional differences.

Consistent with the unadjusted comparisons, financial funds hold significantly higher ESG-rated stocks than benchmark funds within the same quarter. Impact funds hold stocks with significantly lower E ratings and higher S ratings than financial funds, although the difference becomes statistically insignificant when standard errors are clustered

These differences reflect both sector allocation and within-sector selection. As shown in Figure 3, impact funds overweight sectors with lower average E ratings, such as utilities and manufacturing, whereas financial funds tilt toward higher-rated sectors such as information and finance. Notably, utilities tend to have relatively high S ratings, whereas information and finance score lower on that dimension. Thus, differences in sector allocations may explain some of the observed variation in fund-level E and S scores. Panel B formally decomposes these effects.

To separate sector allocation from within-sector stock selection, Panel B reports two decompositions. Columns (1)–(3) assign each stock its sector-average ESG score to isolate the contribution of sector tilts. Columns (4)–(6) use sector-adjusted scores (relative to 6-digit NAICS averages) to capture within-sector selection.

Financial funds overweight sectors with higher average ESG ratings and select firms with higher ESG scores relative to sector peers. In contrast, impact funds overweight sectors with lower E ratings and allocate to firms with weaker ESG performance within sectors. These patterns indicate that financial funds tilt toward already "green" firms and sectors, whereas impact funds target firms and sectors with greater scope for improvement.

Overall, the results reinforce the heterogeneity within sustainable funds: financial funds concentrate in high-ESG firms and sectors, while impact funds allocate capital toward weaker ESG performers, consistent with engagement-oriented objectives.

We further decompose the environmental score (Internet Appendix Table IA4). Impact funds hold stocks with higher Pollution & Waste and Environmental Opportunities scores, consistent with investment in firms whose core activities contribute to environmental solutions. At the same time, impact funds hold firms with lower Natural Capital scores relative to financial funds. This pattern suggests that impact funds allocate capital both to environmentally oriented business models and to firms with weaker environmental performance, consistent with an engagement-based strategy aimed at improvement.

3.3 Carbon Intensity and Environmental Exposure

The environmental score decomposition suggests distinct approaches: financial funds concentrate in firms with stronger environmental ratings, whereas impact funds invest in both environmental solution providers and firms with greater scope for improvement. Because MSCI ratings provide a relative assessment and may not fully capture firms' environmental footprints, we next examine portfolio carbon intensity.

Because financial funds prioritize hedging climate-related risks, we expect them to hold portfolios with lower carbon intensity than benchmark funds. In contrast, impact funds aim to generate positive externalities by investing in higher emission firms and seeking to reduce emissions through engagement.

Carbon intensity is measured as Scope 1&2 and Scope 1&2&3 emissions in metric tons per million dollars of revenue. We also construct sector-average carbon intensity (based on 6-digit NAICS industries) and within-sector-ranked intensity measures.

Panel A of Table 4 reports mean carbon intensity by fund type (columns (1)–(4)) and regression-based differences between fund types, controlling for year–quarter fixed effects (columns (5)–(7)). Columns (5)–(7) show that financial funds hold portfolios with significantly lower carbon intensity than benchmark funds, while impact funds hold portfolios with significantly higher carbon intensity than both benchmark and financial funds. These findings are robust to clustering standard errors at the fund level.

Panel B decomposes carbon intensity into sector allocation and within-sector selection components, mirroring the approach in Table 3. Financial funds overweight lower-emission sectors but do not differ significantly from benchmark funds within sectors. In contrast, impact funds overweight higher-emission sectors while selecting firms with lower emissions intensity relative to sector peers.

These patterns suggest that financial funds reduce portfolio carbon exposure primarily through sector tilts, whereas impact funds allocate capital to carbon-intensive industries but focus on firms that are comparatively better positioned within those sectors—consistent with targeting transition-relevant firms rather than simply holding the highest emitters.

4. Sustainable Funds Performance and Actions

4.1 Performance and Flow–Performance Sensitivity by Sustainable Fund Type

This section examines whether mutual funds' stated sustainable investment strategies are associated with differences in financial performance and flow–performance sensitivity. Prior literature finds that socially responsible funds tend to exhibit lower flow–performance sensitivity than conventional funds, consistent with the idea that investors with nonpecuniary preferences are less likely to withdraw capital after underperformance (Bialkowski and Starks (2015); Renneboog, Ter Horst, and Zhang (2011); Bollen (2007)). Motivated by this, we formulate two conjectures. First, moral and impact funds may earn lower alphas than benchmark funds, while financial funds are expected to deliver performance comparable to benchmarks. Second, financial funds should exhibit similar flow–performance sensitivity as benchmark funds, whereas moral and impact funds should exhibit lower sensitivity.

Importantly, both conjectures are joint hypotheses: they require not only that fund managers pursue their stated sustainable investing strategies, but also that investors sort into funds aligned with their preferences. If investors fail to distinguish between fund types—due to search frictions or a lack of transparency—then performance and flow patterns may not vary across categories. For the performance and flow analyses in Tables 5 and 6, we restrict the sample to U.S. domestic equity funds to improve comparability across fund types.

Table 5 presents the performance results. Panel A reports time-series regressions of value-weighted net-of-fee monthly excess returns using the CAPM and Fama–French three-factor models. Panel B compares each sustainable fund type with benchmark funds. We find that financial and impact funds exhibit significantly lower loadings on HML, indicating a greater tilt toward growth stocks. For financial funds, this tilt may reflect their heavier allocations to the technology sector (see Figure 3). For impact funds, this pattern is consistent with investment in firms associated with environmental opportunities (see Internet Appendix Table IA4). In contrast, moral funds exhibit factor exposures statistically indistinguishable from those of benchmark funds.

None of the sustainable fund types earns an alpha statistically different from benchmark funds. However, the point estimates are economically meaningful: both moral and impact funds display negative and relatively large alphas, whereas financial funds' alphas are closer to zero. These patterns are directionally consistent with our conjecture that moral and impact strategies may involve costs not borne by financial or benchmark strategies.

Table 6 investigates how flows respond to performance. Following Barber, Huang, and Odean (2016), we measure performance using CAPM alpha, which they find has the strongest predictive power for mutual fund flows, and construct it as an exponentially weighted average of prior monthly alphas over an 18-month window. Column (1) compares financial and benchmark funds. The interaction term between Alpha and Fund Type is negative, indicating somewhat weaker flow–performance sensitivity for financial funds, but the difference is not statistically significant. Thus, while financial funds may attract some investors with nonpecuniary motives, the attenuation in flow–performance sensitivity is not estimated precisely enough to support a clear inference.

Column (2) compares moral and benchmark funds. Moral funds exhibit significantly weaker flow–performance sensitivity than benchmark funds. A one-percentage-point decline in alpha reduces flows by 0.824% for benchmark funds, but by only 0.556% (= 0.824 − 0.268) for moral funds. This pattern is consistent with moral funds attracting investors who derive utility from ethical exclusionary screens and are therefore less responsive to performance.

Column (3) compares impact and benchmark funds. As with financial funds, the interaction term is negative, indicating weaker flow–performance sensitivity than among benchmark funds, but the difference is not statistically significant. Thus, although the point estimate suggests some attenuation, we cannot conclude that impact funds are systematically matched with less performance-sensitive investors. One possible explanation for the imprecise impact-fund result is that impact-oriented investors may have more difficulty identifying funds aligned with their preferences. Relative to moral strategies, which often rely on readily observable exclusionary screens, impact strategies tend to involve more complex and multidimensional objectives—such as engagement or thematic investing—that are harder to infer from standard disclosures. These results across all three sustainable fund types are robust to using an alternative three-month decay function and to excluding fixed effects; see Internet Appendix Table IA5.

Taken together, the results indicate that weaker flow–performance sensitivity is most clearly concentrated in moral funds. At the same time, the negative interaction estimates for financial and impact funds suggest that nonpecuniary motives may not be exclusive to moral-fund investors, even if the evidence outside the moral category is imprecise. We view the flow evidence as a relatively demanding and indirect test of fund-type heterogeneity, since it depends not only on how well stated objectives are measured, but also on how accurately investors sort into funds aligned with those objectives. We therefore view the flow evidence as informative, but not decisive, about differences in investor clientele across fund types.

4.2 Proxy Voting on Environmental and Social Shareholder Proposals

We now examine whether funds' voting behavior on shareholder proposals aligns with their stated investment objectives. Proxy voting is a key mechanism through which institutional investors can influence corporate behavior, particularly with respect to environmental and social

(ES) issues (Iliev and Lowry (2015), Brav et al. (2024), Di Giuli et al. (2025)).[23] If impact funds are genuinely motivated to generate positive externalities, we would expect them to exhibit higher support for ES-related shareholder proposals, even when such proposals may impose costs on firms. In contrast, financial funds are expected to support ES proposals when they align with long-term value creation or mitigate transition, regulatory, or reputational risks. To the extent that financial funds systematically integrate ESG considerations into valuation, they should support ES proposals more frequently than benchmark funds, but less frequently than impact funds. The voting behavior of moral funds is more ambiguous ex ante, given the diverse ethical concerns and norms that may motivate this group.

Following He, Kahraman, and Lowry (2023), we identify ES-related proposals using a dummy variable and calculate, for each fund-quarter, the percentage of such proposals that a given fund supported. We also examine the incidence of opposing, abstaining from, or not casting a vote on ES proposals, enabling a more granular analysis of voting behavior across fund types.[24]

Table 7 reveals marked differences in voting behavior across fund categories. Panel A shows that each type of sustainable fund—financial, moral, and impact—supports ES proposals at substantially higher rates than benchmark funds (40.7%–58.5% vs. 25.1%). This pattern suggests that, regardless of specific motivation, sustainable funds exhibit greater engagement on ES issues. Sustainable funds also have significantly lower rates of abstention and non-voting, indicating a more active stance in proxy governance.

Importantly, substantial heterogeneity exists within the sustainable category. Impact funds show the highest support for ES proposals (58.5%), followed by moral funds (47.0%), and

[23] The literature examines the relative effectiveness of exit versus engagement in achieving social impact. Edmans et al. (2022) study conditional and unconditional exclusion of brown firms, and Broccardo et al. (2024) show that exit is ineffective unless a large share of investors act in a socially responsible manner. Relatedly, Dyck et al. (2019) provide cross-country evidence that institutional investors can improve firms' environmental and social performance and argue that this occurs mainly through engagement with portfolio firms rather than through exit. In this section, we focus on the engagement decision.

[24] We do not exclude sustainable index funds from our sample, as Appel et al. (2016) suggest that passive mutual funds also influence firms' governance. While Heath et al. (2022) find that passive index funds monitor less, passive index sustainable funds are a minority in our sample and inclusion would only bias against finding positive results.

financial funds (40.7%). This gradient aligns with our conceptual framework: impact funds have the strongest incentives to use proxy voting to drive improvements in corporate ESG performance, while financial funds are more selective, backing proposals they deem financially material. Moral funds fall in between, reflecting an orientation toward ethical values rather than impact generation or financial performance per se.

While all moral funds by construction have the largest percentage of sustainable sentences classified as moral, some of them also have secondary sustainability intent that may affect their voting behavior. We thus define pure moral funds as those with only moral sentences (thus no financial or impact sentences) and compare them to other moral funds. We find that pure moral funds' support for ES proposals is nearly identical to that of financial funds. In contrast, the higher support for environmental and social (ES) proposals among moral funds is mainly driven by moral funds with also some impact sentences. The results are reported in the Internet Appendix Table IA6.

Panel B presents regression results controlling for year–quarter fixed effects. In Panel B-1, we find that financial funds are 11.8 percentage points more likely to vote in favor of ES proposals than benchmark funds, 9.6 percentage points less likely to vote against, and 2.6 percentage points less likely to abstain. Panel B-2 shows that, relative to other sustainable funds, impact funds are an additional 19.8 percentage points more likely to vote in favor of ES proposals and 20.8 percentage points less likely to vote against. These effects are economically large and statistically significant.

Next, we focus on outcome-focused E&S proposals and examine whether impact funds vote differently on these proposals. If impact funds are truly impact-driven, they should be especially supportive of outcome-focused proposals relative to financial or moral funds. In contrast, for other E&S proposals (e.g., those that advocate for financially material disclosures), the behavioral differences among the three sustainability fund types are a priori less clear.[25] Panel C reports regression results estimated at the fund–proposal level, where the dependent

[25] We thank Nadya Malenko for suggesting this analysis.

variable is an indicator equal to one if the fund voted Yes on the proposal. We define a proposal as Outcome Proposal if it falls into an ISS category whose proposal descriptions (*AgendaGeneralDesc* and *ItemDesc*) are not related to ESG reporting.[26]

Column (1) includes all sustainable funds, while column (2) compares impact funds with benchmark funds. In both specifications, the interaction between the Impact Fund and Outcome Proposal indicators is positive and statistically significant, indicating that impact funds are more likely to support outcome-focused E&S proposals than disclosure-focused ones. Among all sustainable funds, impact funds are 18.1 percentage points more likely to vote in favor of outcome-focused proposals, increasing to 21.2 percentage points relative to benchmark funds. Outcome-focused proposals receive lower average support, consistent with higher implementation costs. Overall, the results suggest that impact funds disproportionately support proposals aimed at achieving real outcomes despite their lower average support. Taken together, the voting patterns documented in Table 7 provide strong evidence that funds act in ways that reflect their stated objectives. Impact funds' elevated support for ES proposals is consistent with the use of proxy voting as a deliberate tool for engagement and influence. Financial funds, while more supportive than conventional peers, adopt a more measured approach—likely weighing financial materiality. Moral funds exhibit intermediate behavior, consistent with a values-based orientation that varies across ethical dimensions.

These findings complement our earlier results on portfolio composition. Impact funds not only invest in firms with the potential for ESG improvement but also use governance mechanisms[27] to encourage such improvements. The alignment between stated objectives and observed actions—across both asset allocation and voting behavior—offers new insight into how different sustainable investment motivations are operationalized in practice.

---

[26] This definition consists of proposals with the following ISSAgendaItem IDs: S0224, S0411, S0416, S0703, S0732, and S0745. These proposals seek to eliminate corporate involvement in harmful products and fossil fuels and to implement policies that promote employment welfare and human rights standards. The detailed descriptions are provided in Internet Appendix Section 5.

[27] See Michaely, Rubio, and Yi (2023) on voting as a communication channel between institutional investors and firms.

5. Emission Outcomes

5.1 E and S Performance Improvements During Holding Periods

5.1.1 ESG Ratings Analysis

A critical test of impact funds' effectiveness is whether they actually succeed in improving the environmental and social performance of their portfolio companies. If impact funds genuinely select firms with poor E and S performance but improvement potential—as suggested by our findings in Sections 3.2 and 3.3—and use proxy voting as one observable channel of engagement (as shown in Section 4.2), we should observe greater improvements in ESG outcomes for companies held by impact funds compared to those held by other fund types.

To evaluate this, we examine changes in two sets of outcome metrics: ESG ratings and carbon emission intensity. For each new stock holding reported by a fund in month t, we track the change in ESG, E, and S ratings (as well as emission intensity) at t+12, t+18, and t+24. Stocks are included only if they remain in the fund portfolio at each respective time point, [28] which introduces survivor bias—funds may sell firms whose E and S performance deteriorates, selectively retaining firms that improve. However, as long as this culling behavior is similar across fund types, observed differences in improvement rates among the remaining holdings can still be informative about fund managers' influence.

Panel A of Table 8 presents average changes in ESG, E, and S ratings across benchmark, financial, moral, and impact funds. We find that, in general, ESG ratings improve over time among stocks that remain in fund portfolios. Comparing benchmark and financial funds, we do not observe consistent differences in the extent of ESG rating improvements: for example,

[28] In cases where a fund's classification changes during a stock's holding period—due to updates in the fund's prospectus language—we treat the investment as exited in the month when the fund is no longer classified under the type it was assigned at the time of the stock's purchase. As a robustness check, we allow the holding period to continue beyond the classification "switch" as long as the fund retains at least one sentence associated with the original category. For example, if a fund initially classified as an impact fund later becomes classified as a moral fund but still contains at least one sentence classified as impact (i.e., becomes a hybrid moral–impact fund), we continue to treat the investment as ongoing until the stock is sold. Our main results are robust to this alternative specification.

financial fund holdings show smaller improvements in E ratings but somewhat larger improvements in S ratings relative to benchmark funds. Results for moral funds are similarly mixed.

By contrast, impact funds exhibit consistently stronger improvements in both E and S ratings relative to benchmark funds. Although ESG ratings may not perfectly capture real-world improvements, this evidence suggests that impact fund investments are associated with positive ESG performance trajectories at the portfolio firm level.

To examine more direct environmental outcomes, we turn next to changes in portfolio companies' carbon emission intensity.

5.1.2 Carbon Emission Improvements During Holding Periods

In this section we examine changes in emission intensity at the investment holdings level. We define investment holding periods analogously to the previous section. The sample sizes are different because MSCI and Trucost coverage of firms differ (Trucost covers more firms than MSCI). We measure both changes in Scope 1 & 2 and Scope 1, 2 & 3 emission intensity from the purchase month t to t+12, t+18 and t+24 for benchmark, financial, moral and impact funds.

Panel B of Table 8 reports the average change in carbon emission intensity during investment holding periods across fund types. On average, we observe a decline in emission intensity for portfolio companies held by all fund categories. This pattern likely reflects both the survivor bias discussed in the prior ESG analysis—where firms that worsen on emissions may be dropped from portfolios—and a broader secular trend of declining corporate emission intensity during the sample period, as firms adopted more carbon-efficient technologies.

Comparing sustainable and benchmark funds, we find that sustainable fund holdings generally exhibit greater reductions in emission intensity. However, the magnitude of emission intensity decline is not significantly different between financial and impact funds. In contrast, moral fund holdings show more modest reductions. These findings suggest that while sustainable funds as a group are associated with greater carbon intensity improvements than benchmark funds, there is heterogeneity in performance across sustainable fund types.

5.2 Emission Intensity Changes: Difference-in-Differences Analysis

A limitation of the previous analysis is that it does not account for heterogeneity in the types of portfolio companies held by different funds. Reducing emission intensity for an airline company, for example, is fundamentally different from doing so for a cement producer. Additionally, overlap in fund ownership may confound attribution: if a firm purchased by a financial fund was also held—currently or previously—by an impact fund, any observed reduction in emissions could be mistakenly attributed to the financial fund. Additionally, we observe that firms held by all sustainable type funds exhibit reductions in emission intensity. This decline in emissions may arise either because sustainable funds select firms that are already on a trajectory of declining emission intensity (a selection effect), or because fund ownership influences firm management and leads to reductions in emissions (a treatment effect). To disentangle these two effects, we (i) implement a stacked difference-in-differences (DiD) framework and (ii) estimate the dynamic effects of sustainable fund ownership. We adopt this specific DiD implementation for two reasons.

First, the recent DiD literature highlights problems that arise when treated firms are used as controls under staggered treatment timing (Callaway and Sant'Anna, 2021; Baker, Larcker, and Wang, 2022). In our setting, where firms are held by different impact funds at different times, this concern naturally applies. Therefore, for each treated firm, we identify a matched control firm based on *ROA* and *log(total assets)* within the same 2-digit NAICS sector. Control firms are further required to have received no sustainable fund investment between t–24 and t+24.We stack the treatment firm and matched control firms as our sample of analysis. This way, the treated firm is never used as control firms and alleviate the staggered DID concern. Moreover, as our control firms and treatment firms are from the same sector, this allows us to account for heterogeneity in the types of portfolio companies held by different funds and be able to compare emission within sector.

Second, treated firms may be held by sustainable funds multiple times during the sample period, generating multiple treatment events. To accommodate settings with multiple treatments in close succession, we follow Sandler and Sandler (2014) and allow multiple event-time dummies to be activated simultaneously. This approach allows us to compare treated and control firms prior to treatment and to trace the dynamic evolution of emission reductions, helping to distinguish between selection and treatment effects. Specifically, we estimate a dynamic DiD model as follows:

$$y_{i,t} = \sum_{k=-18}^{18} Holding_{i,t}^{k} + FEs + \varepsilon_{i,t}$$

where $Holding_{i,t}^{k}$ is a dummy variable equal to one for firm i at calendar time t if t is k months before/after the first month in which firm i becomes held by a fund of a given sustainability type for $k \in [-18,18]$.[29] We define a fund's holding session as the uninterrupted period during which a stock appears in the fund's portfolio Holding sessions lasting fewer than six months are excluded from the analysis. The outcome variable is Scope 1&2 and Scope 1&2&3 carbon emission intensity. All regressions include firm and year-month fixed effects, and standard errors are clustered at the firm level. We drop the constant term and retain the variable at $Holding_{i,t}^{k}$ at k = 0, so that we can directly compare treated and control firms at the time of treatment.

Figure 4 plots the estimated dynamic coefficients for Scope 1&2 emission intensity. For impact funds (Panel A of Figure 4), we find no evidence of pre-trend violations: treated firms and matched controls follow parallel trajectories prior to investment. Starting around month +10, treated firms exhibit significantly lower emission intensity, with effects lasting approximately six months and remaining negative thereafter. These results are consistent with the hypothesis that

[29] Suppose firm i is held by two impact funds that begin holding in January 2018 and January 2019, respectively. Event-time indicators are defined relative to each fund's initiation date and shift forward mechanically over calendar time. For example, 18 months before the January 2018 initiation, the corresponding −18 dummy equals one; as time advances, the active dummy moves toward 0 and then to positive event time. When holding windows overlap, multiple event-time dummies can be activated simultaneously (e.g., one reflecting Fund 1 at event time −6 and another reflecting Fund 2 at event time −18). Event-time indicators are defined over the window [−18, +18] and equal zero outside this range. If two funds of the same type begin holding firm i at the same time, we set the indicator equal to one. Internet Appendix Figure IA1 reports robustness results in which such indicators are replaced by the number of simultaneous holding events sharing the same entry date.

impact fund engagement contributes to meaningful reductions in emissions, with a lag that reflects the time required for operational change.

The timing of emission reductions among firms held by impact funds helps disentangle whether these reductions reflect impact funds selecting firms that are already on a trajectory of declining emission intensity (a selection effect) or whether fund ownership influences firm management and leads to reductions in emissions (a treatment effect). If impact funds endogenously select firms that are already on a path of declining emission intensity, treated firms would exhibit lower emissions than control firms even prior to being held by impact funds, along with a downward pre-trend. Instead, we find that the significant decline in emissions occurs only after firms enter impact fund portfolios, with no evidence of differential pre-trends. These results suggest that being held by impact funds has a treatment effect, in that impact fund ownership reduces portfolio firms' emission intensity.

While we cannot fully rule out the possibility that impact funds happen to invest in firms at the exact moment their emissions begin to decline, this scenario is not likely. It would require funds to accurately anticipate the precise timing of emission reductions, an event typically driven by firm-specific and largely unobservable operational, technological, or regulatory factors. Therefore, it is highly unlikely that the observed reduction is driven by impact funds selecting firms that are just on the verge of declining emission intensity.

Our definition of impact funds does not distinguish between environmentally focused and socially focused funds. To strengthen the analysis, we restrict the sample to firms held by environmentally focused impact funds, defined as those in the top quintile of support for environmental-related shareholder proposals. The results are presented in Panel B of Figure 4. Consistent with Panel A, we find no evidence of pre-trend violations. Beginning around month +7, treated firms exhibit significantly lower emission intensity relative to control firms, with effects persisting for the remaining twelve months during the analysis period. These findings indicate that the reduction in emission intensity is economically significant. The results for Scope 1&2&3 are reported in the Internet Appendix Table IA7. We find robust results, with no evidence of pre-trends, and treated firms exhibiting significantly lower emission intensity after treatment. The stronger result for environmentally focused impact funds provides corroborating evidence that the fund's sustainability intent is associated with real-world outcomes.

For financial funds (Panel C of Figure 4), we observe anomalous pre-trend patterns: treated firms show significantly lower emission intensity than controls beginning as early as –12 and persisting through approximately +8. The coefficients gradually revert to zero after that point. These pre-treatment patterns suggest that observed emission reductions may stem from prior or concurrent ownership by impact funds. Given that impact fund effects appear to begin around +10, lagged influence from earlier impact fund holdings may be confounding the financial fund results.

To test this, we impose a stricter restriction: we exclude any financial fund investment in firms that were held by impact funds between –18 and +18. Results from this cleaner specification (Panel D of Figure 4) show that pre-trend anomalies largely disappear, except for modest significance between –5 and –2. Crucially, we find no treatment effect after financial fund investment. This suggests that financial fund ownership alone is not associated with emission reductions at portfolio firms.

For moral funds (Panel E and Figure 4), we find weak evidence of pre-trends but no treatment effects. This pattern suggests that moral funds may select firms already on an improving trajectory with respect to emission intensity, but their ownership does not appear to induce further improvement.

Taken together, these results corroborate and extend our earlier findings. Impact funds select firms with relatively poor initial ESG performance (Sections 3.2 and 3.3), actively engage through proxy voting (Section 4.2), and achieve measurable improvements in environmental performance during their holding periods (Section 5.1). The consistency of evidence across portfolio composition, engagement, and outcome dimensions strengthens confidence in our classification methodology and demonstrates that impact funds' actions align with their stated objectives.

The ESG improvement analysis also provides important insights into the real-world consequences of different sustainable investing approaches. Financial funds tend to invest in already “green” companies, and moral funds avoid controversial sectors through exclusionary screens. By contrast, impact funds are uniquely positioned to contribute directly to environmental improvements through their ownership and engagement strategies. Despite managing only 2% of sustainable fund assets—and just 0.08% of total mutual fund assets—impact funds may generate disproportionate positive externalities compared to other sustainable investment strategies. This highlights a key tension in sustainable finance: the fastest-growing categories of ESG investment are not necessarily those with the greatest potential to deliver meaningful environmental and social change. Policymakers, fund managers, and investors must grapple with this disconnect in order to better align capital flows with sustainability outcomes.

6. Conclusion

This paper studies how mutual funds articulate sustainability goals and whether those stated intentions are informative about investment behavior and real-economy outcomes. Using a text-based classification of fund disclosures, we distinguish between financial, moral, and impact-oriented sustainability intent. The classification is validated at the sentence level and serves as an organizing variable for the empirical analysis.

We document that most funds using ESG language articulate sustainability primarily in financial terms, while relatively few express a distinct impact-oriented mandate. This descriptive pattern highlights the importance of distinguishing among different sustainability objectives that are often grouped together under a single ESG label. We show that sustainability intent predicts systematic differences in portfolio composition, investor behavior, and stewardship. In particular, moral funds exhibit lower flow-performance sensitivity, while impact-oriented funds are more likely to support outcome-oriented environmental and social shareholder proposals.

Turning to real-economy outcomes, we find that firms added to the portfolios of impact-oriented funds subsequently reduce carbon intensity relative to matched firms, while no comparable pattern is observed for firms added by financial or moral funds. These results suggest that distinctions in stated sustainability intent are informative about downstream firm behavior. At the same time, the absence of post-investment emissions changes for financial and moral funds underscores the need to separate sustainability intent from broader ESG labeling when evaluating claims about real-world impact.

Our findings also help clarify how sustainability-oriented funds are interpreted in public and academic discussions. Using fund disclosures, we find that the majority of assets in sustainability-labeled mutual funds—including those managed by large asset managers—are allocated to funds with financially motivated sustainability objectives rather than impact-oriented mandates. These funds emphasize the use of ESG information for risk management and valuation, rather than engagement aimed at generating measurable negative externalities reductions. Consistent with this distinction, we find limited evidence that financially motivated funds pursue outcome-oriented strategies at scale. Taken together, these patterns underscore the importance of transparent frameworks for distinguishing among different sustainability objectives.

Several directions for future research emerge from our findings. An open question is why assets are so heavily concentrated in financially motivated sustainability funds, and whether this reflects investor preferences, regulatory constraints, or career incentives faced by fund managers. Future work could also examine how disclosure regimes, such as the EU's Sustainable Finance Disclosure Regulation, affect fund strategies and capital allocation, and whether similar patterns arise in other institutional settings, including pensions, credit markets, and private equity.

References

Abis, Simona, 2022, “Man vs. Machine: Quantitative and Discretionary Equity Management”, unpublished working paper.

Abis, Simona, Andrea M. Buffa, and Meha Sadasivam, 2024. “Different Shades of ESG Funds”, unpublished working paper.

Abis, Simona, and Anton Lines, 2024, “Broken Promises, Competition, and Capital Allocation in the Mutual Fund Industry”, Journal of Financial Economics, 162, 103948.

Akey, P. and Appel, I., 2020. Environmental Externalities of Activism. Unpublished working paper.

Allcott, H., Montanari, G., Ozaltun, B. and Tan, B., 2023. “An economic view of corporate social impact”, forthcoming in the Journal of Finance.

Appel, Ian R., Todd A. Gormley, and Donald B. Keim, 2016. Passive investors, not passive owners. *Journal of Financial Economics*, *121*(1), pp.111-141.

Atta-Darkua, V., Glossner, S., Krueger, P. and Matos, P., “Decarbonizing Institutional Investor Portfolios: Helping to Green the Planet or Just Greening Your Portfolio?”, unpublished working paper.

Avramov, D., Cheng, S., Lioui, A. and Tarelli, A., 2022. Sustainable investing with ESG rating uncertainty. *Journal of financial economics*, *145*(2), pp.642-664.

Baker, A.C., Larcker, D.F. and Wang, C.C., 2022. How much should we trust staggered difference-in-differences estimates?. *Journal of Financial Economics*, *144*(2), 370-395.

Baker, M., M. Egan, and S. Sarkar (2022) *How Do Investors Value ESG?*. Unpublished working paper.

Barber, B. M., X. Huang, and T. Odean, 2016. Which factors matter to investors? Evidence from mutual fund flows. *The Review of Financial Studies*, *29*(10), pp.2600-2642.

Barber, Brad, Adair Morse, and Ayako Yasuda (2021) “Impact Investing” Journal of Financial Economics 139, 162-185

Bauer, R., Ruof, T., Smeets, P., 2021. Get Real! Individuals Prefer More Sustainable Investments. *The Review of Financial Studies*, 34, Issue 8, p. 3976–4043.

Benson, K.L. and Humphrey, J.E., 2008. Socially responsible investment funds: Investor reaction to current and past returns. *Journal of Banking & Finance*, *32*(9), pp.1850-1859.

Berg, F., Kölbel, J.F. and Rigobon, R., 2022. Aggregate confusion: The divergence of ESG ratings. *Review of Finance*, *26*(6), pp.1315-1344.

Berk, Jonathan, and Jules H. van Binsbergen. (2025) "The impact of impact investing." Journal of Financial Economics 164, 103972.

Białkowski, Jedrzej, and Laura T. Starks, 2018. SRI funds: Investor demand, exogenous shocks and ESG profiles. Working Paper.

Bingler, Julia and Kraus, Mathias and Leippold, Markus and Webersinke, Nicolas (2024) "How Cheap Talk in Climate Disclosures relates to Climate Initiatives, Corporate Emissions, and Reputation Risk", *Journal of Banking & Finance*, *164*, 107191.

Bini, P., Cong, L.W., Huang, X. and Jin, L.J., 2026. "Behavioral economics of AI: LLM biases and corrections" Unpublished working paper.

Bollen, Nicolas PB., 2007. Mutual fund attributes and investor behavior. *Journal of financial and quantitative analysis*, *42*(3), pp.683-708.

Bonnefon, J-F, A. Landier, P. Sastry, and D. Thesmar, (2025) "The moral preferences of investors: Experimental evidence", *Journal of Financial Economics* 163 103955.

Brav, A., W. Jiang, T. Li, and J. Pinnington, (2024). Shareholder monitoring through voting: New evidence from proxy contests. *The Review of Financial Studies*, *37*(2), 591-638.

Briscoe-Tran, H., 2026. Do employees have useful information about firms' ESG practices? forthcoming in *Management Science.*

Broccardo, Eleonora, Oliver Hart, and Luigi Zingales, (2022). Exit versus voice. *Journal of Political Economy*, *130*(12), pp.3101-3145.

Callaway, B. and Sant'Anna, P.H., 2021. Difference-in-differences with multiple time periods. *Journal of econometrics*, *225*(2), pp.200-230.

Cao, Sean Shun, Wei Jiang, and Lijun Gillian Lei, 2024. Applied AI for finance and accounting: Alternative data and opportunities. *Pacific-Basin Finance Journal*, *84*, p.102307.

Chen, T., Dong, H. and Lin, C., 2020. Institutional shareholders and corporate social responsibility. *Journal of Financial Economics*, *135*(2), pp.483-504.

Chowdhry, Bhagwan, Shaun William Davies, and Brian Waters (2019) "Investing for Impact" The Review of Financial Studies 32: 864–904.

Christensen, D.M., Serafeim, G. and Sikochi, A., 2022. Why is corporate virtue in the eye of the beholder? The case of ESG ratings. *The Accounting Review*, *97*(1), pp.147-175.

Cohen, L., Gurun, U.G. and Nguyen, Q.H., 2020. *The ESG-innovation disconnect: Evidence from green patenting* (No. w27990). National Bureau of Economic Research.

Couvert, Maxime, 2025, "What is the Impact of Mutual Funds' ESG Preferences on Portfolio Firms?", unpublished working paper.

Dangl, Thomas, Michael Halling, Jin Yu, and Josef Zechner, "Social preferences and corporate investment", forthcoming in Journal or Financial Studies.

Dasgupta, Amil, Jenter, Dirk, Mathews, Richmond D. and Voss, Paul, 2025, "Are Socially Responsible Funds Viable?", unpublished working paper.

Di Giuli, Alberta, Alexandre Garel, Roni Michaely, and Arthur Romec, 2025. Climate change and mutual fund voting on climate proposals. *Management Science*.

Dimson, E., Karakaş, O. and Li, X., 2015. Active ownership. *The Review of Financial Studies*, *28*(12), pp.3225-3268.

Dimson, E., Karakaş, O. and Li, X., 2025. Coordinated engagements. *European Corporate Governance Institute–Finance Working Paper*, *721*(6).

Duchin, R., J. Gao, and Q. Xu, (2025) "Sustainability or Greenwashing: Evidence from the Asset Market for Industrial Pollution", *The Journal of Finance* 80, no. 2 (2025): 699-754.

Dyck, A., Lins, K., Roth, L., and Wagner, H. (2019). Do institutional investors drive corporate social responsibility? International evidence. *Journal of financial economics 131*, 693-714.

Edelen, Roger M, 1999. Investor flows and the assessed performance of open-end mutual funds. *Journal of Financial Economics*, *53*(3), pp.439-466.

Edmans, Alex, Tom Gosling, and Dirk Jenter, 2026. Sustainable investing: Evidence from the field. Unpublished working Paper.

Edmans, A., D. Levit, and J. Schneemeier, 2022. Socially responsible divestment. *Centre for Economic Policy Research Working Paper.*

El Ghoul, S. and Karoui, A., 2017. Does corporate social responsibility affect mutual fund performance and flows?. *Journal of Banking & Finance*, *77*, pp.53-63.

Engle, Robert F., Stefano Giglio, Bryan Kelly, Heebum Lee, and Johannes Stroebel, 2020. Hedging climate change news. *The Review of Financial Studies*, *33*(3), pp.1184-1216.

Flammer, C., Toffel, M. and Viswanathan, K., 2021. Shareholder activism and firms' voluntary disclosure of climate change risks. *Strategic Management Journal*, *42*, 1850-1879.

Gantchev, N., Giannetti, M. and Li, R., 2022. Does money talk? Divestitures and corporate environmental and social policies. *Review of Finance*, *26*(6), pp.1469-1508.

Gantchev, N., Giannetti, M. and Li, R., 2024. Sustainability or performance? Ratings and fund managers' incentives. *Journal of Financial Economics*, 155, p.103831.

Gentzkow, Matthew, Bryan Kelly, and Matt Taddy. (2019) "Text as data." Journal of Economic Literature 57, 535-574.

Gibson Brandon, R., Glossner, S., Krueger, P., Matos, P. and Steffen, T., 2022. Do responsible investors invest responsibly?. *Review of Finance*, *26*(6), pp.1389-1432.

Giglio, Stefano, Bryan Kelly, and Johannes Stroebel, 2021. Climate finance. *Annual review of financial economics*, *13*(1), pp.15-36.

Giglio, S., Matteo M., Johannes Stroebel, Zhenhao Tan, Stephen Utkus, and Xiao Xu, 2025, Four facts about ESG beliefs and investor portfolios, *Journal of Financial Economics* 164 (2025).

Goldstein, Itay, Alexandr Kopytov, Lin Shen and Haotian Xiang (2022) "On ESG Investing: Heterogeneous Preferences, Information, and Asset Prices" Unpublished working paper

Green, D., and B. Roth. (2025) "The allocation of socially responsible capital." *The Journal of Finance* 80, no. 2 (2025): 755-781.

Gupta, D., Kopytov, A., & Starmans, J. (2026). The pace of change: Socially responsible investing in private markets. *The Review of Financial Studies*, *39*(1), 30-78.

Hartzmark, S. M. and Sussman, A. B. (2019): Do investors value sustainability? A natural experiment examining ranking and fund flows, Journal of Finance 74, 2789–2837.

Hartzmark, Samuel M., and Kelly Shue. (2023) Counterproductive sustainable investing: The impact elasticity of brown and green firms. Unpublished working Paper.

He, Haibo, and Edwardo A. Garcia. (2009) "Learning from imbalanced data." *IEEE Transactions on knowledge and data engineering* 21, no. 9: 1263-1284.

He, Yazhou Ellen, Bige Kahraman, and Michelle Lowry, (2023) “ES Risks and Shareholder Voice”, *The Review of Financial Studies* 36, 4824-4863.

Heath, D., Macciocchi, D., Michaely, R. and C. Ringgenberg, M., 2023. Does socially responsible investing change firm behavior?. *Review of Finance*, *27*(6), pp.2057-2083.

Heeb, Florian, Julian F. Kölbel, Falko Paetzold and Stefan Zeisberger (2023) “Do Investors Care about Impact?” Review of Financial Studies 36, 1737-1787.

Heeb, Florian, Julian Kölbel, and Camilla Weder, (2025) “Beliefs about the Climate Impact of Green Investing”, unpublished working paper.

Hoepner, A.G., Oikonomou, I., Sautner, Z., Starks, L.T. and Zhou, X.Y., 2024. ESG shareholder engagement and downside risk. *Review of Finance*, *28*(2), pp.483-510.

Huang, A.H., Wang, H. and Yang, Y., 2023. FinBERT: A large language model for extracting information from financial text. *Contemporary Accounting Research*, *40*(2), pp.806-841.

Iliev, Peter, and Michelle Lowry, (2015). Are mutual funds active voters?. *The Review of Financial Studies*, *28*(2), 446-485.

Jiang, Wei, Meng Wang, and Baozhong Yang, 2024. Measuring Sustainability with AI. In *Artificial Intelligence, Finance, and Sustainability: Economic, Ecological, and Ethical Implications* (pp. 33-57). Cham: Springer Nature Switzerland.

Jin, Dunhong, Roni Michaely, and Menghan Wang, 2024. How green is green? Anatomy of ESG funds’ selection. Working Paper.

Kim, S. and Yoon, A., 2023. Analyzing active fund managers’ commitment to ESG: Evidence from the United Nations Principles for Responsible Investment. *Management science*, *69*(2), pp.741-758.

Krueger, Philipp, Zacharias Sautner, and Laura T. Starks, 2020, The importance of climate risk for institutional investors, *Review of Financial Studies* 33, 1067–1111.

Landier, Augustin and Lovo, Stefano (2025) “Socially Responsible Finance: How to Optimize Impact” *The Review of Financial Studies* 38, no. 4 (2025): 1211-1258.

Lee, Matthew, Arzi Adbi, and Jasjit Singh (2020) "Categorical cognition and outcome efficiency in impact investing decisions” Strategic Management Journal 41:86–107

Lewellen, J. and Lewellen, K., 2022. Institutional investors and corporate governance: The incentive to be engaged. *The Journal of Finance*, *77*(1), pp.213-264.
Li, J. and Wu, D., 2020. Do corporate social responsibility engagements lead to real environmental, social, and governance impact?. *Management Science*, *66*(6), 2564-2588.
Liang, H., Sun, L. and Teo, M., 2022. Responsible hedge funds. *Review of Finance*, *26*(6), pp.1585-1633.
Lindsey, L.A., Pruitt, S. and Schiller, C., 2024. The cost of ESG investing. *Unpublished working paper.*
Lowry, M., Wang, P. and Wei, K.D., 2025. Are all ESG funds created equal? Only some funds are committed, forthcoming in the Review of Financial Studes.
Michaely, R., Ordonez-Calafi, G. and Rubio, S., 2024. Mutual funds' strategic voting on environmental and social issues. *Review of Finance*, *28*(5), pp.1575-1610.
Michaely, R., S. Rubio, and I. Yi (2023) "Voting Rationales", Unpublished working paper
Morgan, John and Justin Tumlinson (2019) "Corporate Provision of Public Goods" Management Science 65, 4489-4504.
Naaraayanan, S.L., Sachdeva, K. and Sharma, V., 2021. The real effects of environmental activist investing. *European Corporate Governance Institute–Finance Working Paper*.
Oehmke, Martin and Opp, Marcus M. (2025) "A Theory of Socially Responsible Investment" Review of Economic Studies 92, 1193-1225.
O'Hara, M. and Streltsov, A., 2024. Keeping the Faith (and the Returns): An AI Approach to Values-based Investing. *Unpublished Working Paper*.
Parise, G., and M. Rubin, 2025. Green window dressing. forthcoming in the *Journal of Finance.*
Pástor, L., R. Stambaugh, and L. Taylor (2021) "Sustainable Investing in Equilibrium" Journal of Financial Economics, 142(2), 550-571
Pástor, L., R. Stambaugh, and L. Taylor, 2022. Dissecting green returns. *Journal of financial economics*, *146*(2), pp.403-424.
Pedersen, Lasse Heje, Shaun Fitzgibbons, and Lukasz Pomorski (2021) "Responsible Investing: The ESG-efficient Frontier" Journal of Financial Economics 142(2) 572–597.
Raghunandan, A. and Rajgopal, S., 2022. Do ESG funds make stakeholder-friendly investments?. *Review of Accounting Studies*, *27*(3), pp.822-863.
Rajan, R. and Ramella, Pietro and Zingales, Luigi (2022) "What Purpose Do Corporations Purport? Evidence from Letters to Shareholders", Unpublished working paper.
Rajgopal, Shiva, Anup Srivastava, and Rong Zhao, 2025. Economic substance behind Texas political anti-ESG sanctions. *Management Science*.

Renneboog, Luc, Jenke Ter Horst, and Chendi Zhang., 2011. Is ethical money financially smart? Nonfinancial attributes and money flows of socially responsible investment funds. *Journal of Financial Intermediation*, *20*(4), pp.562-588.

Riedl, A. and Smeets, P. (2017), Why Do Investors Hold Socially Responsible Mutual Funds?. The Journal of Finance, 72: 2505-2550.

Sandler, Danielle H., and Ryan Sandler, 2014. Multiple event studies in public finance and labor economics: A simulation study with applications. *Journal of Economic and Social Measurement*, *39*(1-2), pp.31-57.

Sautner, Z., Van Lent, L., Vilkov, G. and Zhang, R., 2023. Firm-level climate change exposure. *The Journal of Finance*, *78*(3), pp.1449-1498.

Shanker, H. (2025) "Ethical Capital, Coordination, and the Correction of Externalities", unpublished working paper.

Starks, Laura (2023) "Presidential Address: Sustainable Finance and ESG Issues—*Value* versus *Values"* Journal of Finance 78, 1837-1872.

Sulaeman, Johan, and Qiaozhi Ye. (2023) "Who Do You Vote For? Same-Race Voting Preferences in Director Elections." Unpublished Working Paper.

U.S. Securities and Exchange Commission. (2016) "How to Read a Mutual Fund Prospectus" from the SEC website.

Appendix

A1: Keywords used to identify ESG-Sentences.
'esg','environment','social','governance','sustainable','sustainability','abortion','lgbt','gay','lesbian','tobacco','gambling','alcohol','pornography','gun','energy','fossil','fuel','green','impact','responsible','clean','minority','minorities','poverty','girl','girls','male','female','fair','maternity','paternity','equal','equality','discrmination','non-discrimination', 'sexual', 'harassment', 'safety', 'diversity', 'civic', 'trafficking','ethics','gender','race','ethnicity', 'climate', 'renewable', 'energy', 'vote', 'voting','proxy', 'transform', 'transformation', 'dialogue', 'engage', 'engagement', 'transition'

A2: Variable Definitions

Variable Name	Definition
Financial$_{it}$	equal to 1 if the most recent available prospectus for fund i as of quarter t contains the highest percentage of Financial sentences among all BERT-classified sustainable sentences (Financial, Moral, and Impact).
Moral$_{it}$	equal to 1 if the most recent available prospectus for fund i as of quarter t contains the highest percentage of Moral sentences among all BERT-classified sustainable sentences (Financial, Moral, and Impact).
Impact$_{it}$	equal to 1 if the most recent available prospectus for fund i as of quarter t contains the highest percentage of Impact sentences among all BERT-classified sustainable sentences (Financial, Moral, and Impact).
Pure Financial$_{it}$	equal to 1 if the most recent available prospectus for fund *i* contains at least one sentence classified as Financial, and no sentences classified as Moral or Impact.
Pure Moral$_{it}$	equal to 1 if the most recent available prospectus for fund *i* contains at least one sentence classified as Moral, and no sentences classified as Financial or Impact
Pure Impact$_{it}$	equal to 1 if the most recent available prospectus for fund *i* contains at least one sentence classified as Impact, and no sentences classified as Financial or Moral
Environmental Impact$_{it}$	equal to 1 if the fund is an impact fund and is in the top quintile based on its support for environmental-related shareholder proposals.
Holding-quarters$_{it}$	the average number of quarters for which the stocks in the portfolio have been held by fund *i* as of quarter *t*.
Retail Expense Ratio$_{it}$	the average expense ratio for retail funds classes for fund *i* as of quarter *t*.
ESG Score$_{jt}$	the average of the evaluated MSCI Environmental and Social key issues scores that firm *j* received from MSCI in quarter *t*
E Score$_{jt}$	the average of the evaluated MSCI Environmental key issues scores that firm *j* received from MSCI in quarter *t*
S Score$_{jt}$	the average of the evaluated MSCI Social key issues scores that firm *j* received from MSCI in quarter *t*

Sector-Average-ESG Score$_{it}$	assign the quarter-sector average *ESG Score*$_{jt}$ at the 6-digit NAICS level to each stock. $\frac{\sum_{for\ all\ k\ in\ the\ industry\ of\ firm\ j}(ESG\ Score_{kt} * MarketValue_{kt})}{\sum_{\text{for all k in the industry of firm j}}(\text{MarketValue}_{\text{kt}})}$
Sector-Average-E Score$_{jt}$	assign the quarter-sector average *E Score*$_{jt}$ at the 6-digit NAICS level to each stock. $\frac{\sum_{for\ all\ k\ in\ the\ industry\ of\ firm\ j}(E\ Score_{kt} * MarketValue_{kt})}{\sum_{\text{for all k in the industry of firm j}}(\text{MarketValue}_{\text{kt}})}$
Sector-Average-S Score$_{jt}$	assign the quarter-sector average *S Score*$_{jt}$ at the 6-digit NAICS level to each stock. $\frac{\sum_{for\ all\ k\ in\ the\ industry\ of\ firm\ j}(S\ Score_{kt} * MarketValue_{kt})}{\sum_{\text{for all k in the industry of firm j}}(\text{MarketValue}_{\text{kt}})}$
Sector-Adjusted-ESG Score$_{jt}$	*ESG Score*$_{jt}$ minus *Sector-Average-ESG Score*$_{it}$.
Sector-Adjusted-E Score$_{jt}$	E Score$_{\text{jt}}$ minus *Sector-Average-E Score*$_{jt}$.
Sector-Adjusted-S Score$_{jt}$	*S Score*$_{jt}$ minus *Sector-Average-S Score*$_{jt}$.
Scope 1&2(&3) Carbon Emission Intensity$_{jt}$	the Scope 1&2(&3) emissions, measured in metric tons per million dollars of revenue, that firm *j* reports at the fiscal year-end immediately following quarter t
Sector-Average Scope 1&2(&3) $_{jt}$	assigning the quarterly sector-level Scope 1&2(&3) emission intensity at the 6-digit NAICS level to each stock $\frac{\sum_{for\ all\ k\ in\ the\ industry\ of\ firm\ j}(Emission * MarketValue_{kt})}{\sum_{\text{for all k in the industry of firm j}}(\text{MarketValue}_{\text{kt}})}$
Within-Sector-ranked Scope 1&2(&3) $_{jt}$	the decile rank of firm *j*'s Scope 1&2(&3) emission intensity within the 6-digit NAICS code.
12 (18, or 24) -Month ESG (E, or S) Score Change	Change in MSCI ESG (E, or S) ratings of firm *j* held by fund *i* between the time of initial investment and 12 (18, or 24) months post-investment
12 (18, or 24) -Month Scope 1&2 Intensity Change	Change in Scope 1&2(&3) Intensity of firm *j* held by fund *i* between the time of initial investment and 12 (18, or 24) months post-investment. A firm's carbon emission at year-month *t* is calculated as the weighted average of the last available and the next available (annually measured) carbon intensity. The weight on the last available carbon intensity equals 12 minus the difference between the current month and the reporting month divided by 12, with the remaining weight applied to the next available carbon intensity.
ES "For" Vote$_{it}$	the percentage of Environmental and Social shareholder proposals received by firms held by fund *i* that fund *i* voted for in quarter *t*
ES "Against" Vote$_{it}$	the percentage of Environmental and Social shareholder proposals received by firms held by fund *i* that fund *i* voted against in quarter *t*
ES "Abstain" Vote$_{it}$	the percentage of Environmental and Social shareholder proposals received by firms held by fund *i* that fund *i* abstained from voting in quarter *t*
ES "Do Not Vote"$_{it}$	the percentage of Environmental and Social shareholder proposals received by firms held by fund *i* that fund *i* did not cast a vote in quarter *t*

Outcome Proposal	equal to 1 if the proposal is in an ISS category whose proposal descriptions by ISS (AgendaGeneralDesc and ItemDesc) are not related to ESG reporting, consisting of ISSAgendaItemIDs S0224, S0411, S0416, S0703, S0732, and S0745.

Table 1: BERT Model Performance

This table reports the performance of the BERT model trained to classify ESG-related sentences. 2,950 sentences in the training sample and 625 in the testing sample are manually labeled into four different categories: Financial, Moral, Impact, or None. The model is trained on the training sample for classification. The trained BERT model is applied to the testing sample to obtain the ESG classification given by the BERT model. Four different model performance measures are calculated to measure the accuracy of BERT classification. Accuracy is the ratio of (true positives + true negatives) divided by the total number of observations (fraction of correct classifications). Precision is the ratio of true positives divided by the sum of true positives and false positives. Recall is the ratio of true positives divided by the sum of true positives and false negatives. *f1* is defined as $[\frac{2}{\frac{1}{Precision}+\frac{1}{Recall}}]$.

	Accuracy	Precision	Recall	f1
Financial	0.9088	0.8467	0.7888	0.8167
Moral	0.9840	0.8500	0.8947	0.8718
Impact	0.9440	0.8276	0.6575	0.7328

Table 2: Summary Statistics

This table reports summary statistics for three categories of sustainable funds from 2014 Q2 to 2023 Q4. The number of funds equals the number of fund-quarter observations divided by 39 (the number of quarters in the sample period). % of equity in the portoflio equals the value of equity holdings divided by the fund's total NAV as of the quarter end. The number of stocks equals the number of common stocks held in the fund portfolio. The fund size is an average among all fund-quarters. The number of funds with retail class is the number of fund-quarter observations for funds with retail class offering(s) divided by 39. The retail expense ratio equals the expense ratio for retail funds. Benchmark funds are CRSP equity non-index funds or ETFs whose prospectuses contain no sentences with sustainability keywords and are never on the Morningstar sustainable fund list during the sample period. A fund is assigned the category with the largest % share of all sustainable sentences.

	Benchmark	Financial	Moral	Impact
# of funds (per quarter)	1524	314	90	46
% of equity in the portfolio	86.0%	88.6%	87.7%	89.8%
# of stocks in the portfolio	138.1	99.3	183.9	90.1
Fund size ($M)	$2,985	$1,226	$769	$327
# of funds with retail class	992	193	51	25
Retail Expense Ratio	1.35%	1.34%	1.27%	1.47%

Table 3: ESG Characteristics of Fund Holdings

This table reports ESG ratings of portfolio holdings by fund type. ESG Score is the average of the evaluated MSCI Environmental and Social Key Issues scores. E Score is the average MSCI Environmental Key Issues scores. S Score is the average MSCI Social Key Issues scores. For each fund-quarter, we calculate the value-weighted ESG ratings for all MSCI-rated stock holdings. In Panel A, columns (1) through (4) report the mean ESG, E, and S scores for benchmark funds, sustainable funds, financial funds, and impact funds, respectively. Columns (5) and (6) report regression results from regressing fund-quarter average ESG, E, and S ratings on the financial fund dummy and the impact fund dummy, controlling for year-quarter fixed effects. In column (5), the sample includes financial funds and benchmark funds, so the coefficient on the financial fund dummy captures the difference between financial and benchmark funds. In column (6), the sample includes impact funds and financial funds, so the coefficient on the impact fund dummy captures the difference between impact and financial funds. In Panel B, columns (1), (2), and (3), the fund's sector-average ESG, E, and S ratings are regressed on the fund indicator variables. To construct the fund's sector-average ratings, we assign the quarter-sector average rating at the 6-digit NAICS level to each stock and calculate the value-weighted fund portfolio rating. In columns (4), (5), and (6), the sector-adjusted ESG, E, and S ratings are regressed on the fund indicator variables. To construct the sector-adjusted ratings, we subtract the quarter-sector average ratings at the 6-digit NAICS level from the raw ratings. Panel B-1's sample includes financial funds and benchmark funds. Panel B-2's sample includes impact funds and financial funds. * , ** , and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Panel A: Average Ratings

	(1) Benchmark	(2) Sustainable	(3) Financial	(4) Impact	(5) Financial - Benchmark	(6) Impact - Financial
ESG Score	5.716	6.056	6.099	5.988	0.160***	-0.011
E Score	6.370	6.729	6.784	6.583	0.173***	-0.104***
S Score	5.059	5.381	5.412	5.392	0.146***	0.085***
Observations	59423	16463	12248	1790		

Panel B: Within-Sector and Cross-Sector Analysis

Panel B-1: Within-Sector and Cross-Sector Financial Funds vs Benchmark Funds

	(1) Sector-Average ESG Score	(2) Sector-Average E Score	(3) Sector-Average S Score	(4) Sector-Adjusted ESG Score	(5) Sector-Adjusted E Score	(6) Sector-Adjusted S Score
Financial Fund	0.058***	0.087***	0.028**	0.102***	0.086***	0.118***
	(3.659)	(3.172)	(2.157)	(4.969)	(3.331)	(6.500)
Observations	71671	71671	71671	71658	71658	71671
Adjusted R2	0.330	0.200	0.364	0.017	0.010	0.027

Panel B-2: Within-Sector and Cross-Sector Impact Funds vs Financial Funds

	(1) Sector-Average ESG Score	(2) Sector-Average E Score	(3) Sector-Average S Score	(4) Sector-Adjusted ESG Score	(5) Sector-Adjusted E Score	(6) Sector-Adjusted S Score
Impact Fund	0.081**	-0.115*	0.277***	-0.093*	0.009	-0.193***
	(2.143)	(-1.675)	(7.489)	(-1.859)	(0.130)	(-3.183)
Observations	13812	13812	13812	13808	13808	13812
Adjusted R2	0.324	0.154	0.358	0.008	0.003	0.028

Table 4: Carbon Emission Intensity of Fund Portfolios

This table reports the carbon emission intensity of the funds' portfolio companies. Carbon Intensity is defined as Scope 1&2, and Scope 1&2&3 Carbon Emission Intensity, measured as metric tons of emissions per million dollars of revenue. In Panel A, columns (1) through (4) report the mean Scope 1&2 Carbon Emission Intensity and Scope 1&2&3 Carbon Emission Intensity for benchmark funds, sustainable funds, financial funds, and impact funds, respectively. Columns (5) through (7) report regression results from regressing Scope 1&2, and Scope 1&2&3 Carbon Emission Intensity on the financial fund dummy and the impact fund dummy, controlling for year-quarter fixed effects. In column (5), the sample includes financial funds and benchmark funds, so the coefficient on the financial fund dummy captures the difference between financial and benchmark funds. In column (6), the sample includes impact funds and benchmark funds, so the coefficient on the impact fund dummy captures the difference between impact and benchmark funds. In column (7), the sample includes impact funds and financial funds, so the coefficient on the impact fund dummy captures the difference between impact and financial funds. Panel B presents regressions of sector-average (columns (1) and (2)) and within-sector-ranked (columns (3) and (4)) Scope 1&2 and Scope 1&2&3 carbon emission intensity on fund indicator variables. A fund's sector-average Carbon Emission Intensity is measured by assigning the quarterly sector-level emission intensity at the 6-digit NAICS level to each stock and then calculating the value-weighted carbon emission intensity. Within-sector-ranked Carbon Emission Intensity is the decile rank of Carbon Emission Intensity within the 6-digit NAICS code. Panel B-1 includes financial funds and benchmark funds, while Panel B-2 includes impact funds and financial funds. Regressions include year-quarter fixed effects, with standard errors clustered at the fund level where indicated. * , ** , and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Panel A: Average Carbon Intensity

	(1) Benchmark	(2) Sustainable	(3) Financial	(4) Impact	(5) Financial - Benchmark	(6) Impact - Benchmark	(7) Impact - Financial
Carbon Intensity Scope 1&2	161.9	132.7	126.4	182.1	-20.3***	28.8***	50.6***
Carbon Intensity Scope 1&2&3	285.4	247.5	235.7	321.4	-22.5***	51.1***	75.6***
Observations	60191	16757	12513	1809			

Panel B: Within-Sector and Cross-Sector Analysis

Panel B-1: Within-Sector and Cross-Sector Financial Funds vs Benchmark Funds

	(1) Sector-Average Scope 1&2	(2) Sector-Average Scope 1&2&3	(3) Within-sector-ranked Scope 1&2	(4) Within-sector-ranked Scope 1&2&3
Financial Fund	-15.534***	-17.702**	0.043	0.039
	(-2.784)	(-2.556)	(1.039)	(0.990)
Observations	72704	72704	72696	72696
Adjusted R^2	0.024	0.047	0.018	0.030

Panel B-2: Within-Sector and Cross-Sector Impact Funds vs Financial Funds

	(1) Sector-Average Scope 1&2	(2) Sector-Average Scope 1&2&3	(3) Within-Sector-ranked Scope 1&2	(4) Within-sector-ranked Scope 1&2&3
Impact Fund	106.620***	136.566***	-0.257**	-0.320***
	(3.146)	(3.764)	(-2.344)	(-2.947)
Observations	14090	14090	14087	14087
Adjusted R^2	0.058	0.091	0.010	0.017

Table 5: Performance and Factor Tilts

This table reports return performance and factor exposures of sustainable funds relative to benchmark funds. The dependent variable is the value-weighted monthly average net-of-fee excess return for each fund type—Financial, Moral, Impact, and Benchmark—over March 2014 to December 2023. For the performance analysis in this table, the sample is restricted to U.S. domestic equity funds in order to improve comparability of returns and factor exposures across fund types. Panel A presents time-series regressions of monthly excess returns for each fund type on the market factor and, in alternate specifications, the Fama–French three factors. The intercept is reported in all specifications. Panel B reports panel regressions at the fund-type-by-month level, comparing each sustainable fund type with the benchmark; interactions between the fund-type indicator and the factor returns capture differences in factor loadings. Columns (1)–(2) compare Financial and Benchmark funds, columns (3)–(4) compare Moral and Benchmark funds, and columns (5)–(6) compare Impact and Benchmark funds. Fund type equals one for Financial funds in columns (1)–(2), Moral funds in columns (3)–(4), and Impact funds in columns (5)–(6).

Panel A: Time Series Regressions

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Benchmark Excess Return		Financial Excess Return		Moral Excess Return		Impact Excess Return	
MktRf	0.973***	0.939***	1.024***	0.993***	1.012***	0.978***	1.030***	0.985***
	(60.611)	(91.329)	(69.404)	(82.420)	(60.174)	(78.309)	(46.170)	(52.168)
SMB		0.170***		0.165***		0.178***		0.251***
		(9.920)		(8.238)		(8.575)		(7.989)
HML		0.110***		0.035**		0.088***		-0.003
		(9.214)		(2.480)		(6.038)		(-0.122)
Constant	-0.201***	-0.133***	-0.173**	-0.121**	-0.269***	-0.202***	-0.251**	-0.182**
	(-2.698)	(-2.898)	(-2.527)	(-2.245)	(-3.439)	(-3.637)	(-2.421)	(-2.158)
Observations	117	117	117	117	117	117	117	117
Adjusted R^2	0.969	0.989	0.976	0.986	0.969	0.984	0.948	0.966

Panel B: Sustainable Funds Compared to Benchmark Funds

	(1)	(2)	(3)	(4)	(5)	(6)
	Financial Excess Return		Moral Excess Return		Impact Excess Return	
MKT	0.973***	0.939***	0.973***	0.939***	0.973***	0.939***
	(63.100)	(83.838)	(59.168)	(82.112)	(50.045)	(61.784)
MKT*Fund Type	0.051**	0.054***	0.040*	0.039**	0.058**	0.046**
	(2.356)	(3.414)	(1.700)	(2.384)	(2.097)	(2.124)
Fund Type	0.028	0.012	-0.068	-0.070	-0.050	-0.049
	(0.275)	(0.173)	(-0.626)	(-0.965)	(-0.390)	(-0.509)
SMB		0.170***		0.170***		0.170***
		(9.106)		(8.919)		(6.711)
HML		0.110***		0.110***		0.110***
		(8.458)		(8.284)		(6.233)
SMB* Fund Type		-0.005		0.008		0.081**
		(-0.173)		(0.313)		(2.271)
HML* Fund Type		-0.076***		-0.023		-0.113***
		(-4.094)		(-1.198)		(-4.515)
Constant	-0.201***	-0.133***	-0.201***	-0.133***	-0.201**	-0.133*
	(-2.809)	(-2.660)	(-2.634)	(-2.606)	(-2.228)	(-1.960)
Observations	234	234	234	234	234	234
Adjusted R2	0.973	0.987	0.969	0.986	0.958	0.977

Table 6: Flow Performance Sensitivity

This table examines the sensitivity of fund flows to performance. For the flow analysis in this table, the sample is restricted to U.S. domestic equity funds in order to improve comparability of performance measures across fund types. For each year-month, alpha is defined as realized excess return minus predicted excess return from the CAPM, where the market beta is estimated over the prior 60 months. The main performance variable, Alpha, is constructed as an exponentially weighted average of monthly alphas from months $t-1$through $t-18$, using a decay parameter of $\lambda = 0.2$. Column (1) compares Financial and Benchmark funds, column (2) compares Moral and Benchmark funds, and column (3) compares Impact and Benchmark funds. Fund Type is an indicator for the sustainable fund type in each specification: it equals one for Financial funds in column (1), Moral funds in column (2), and Impact funds in column (3). Control variables include lagged log total net assets, log fund age, lagged expense ratio, and a load dummy. Standard errors are double clustered by fund and year-month.

	(1) Financial Flow Rate	(2) Moral Flow Rate	(3) Impact Flow Rate
Alpha*Fund Type	-0.276	-0.268**	-0.230
	(-1.524)	(-2.057)	(-0.877)
Alpha	0.831***	0.824***	0.825***
	(9.263)	(9.168)	(9.146)
Fund Type	0.154	0.158	0.190
	(1.093)	(1.172)	(0.862)
Flow Rate at t-19	0.054***	0.054***	0.055***
	(4.677)	(4.229)	(4.187)
Year Month FE	Yes	Yes	Yes
Double Clustering at Fund ×Year Month	Yes	Yes	Yes
Observations	68588	61437	59026
Adjusted R^2	0.007	0.009	0.009

Table 7: Fund Votes on Environmental and Social Shareholder Proposals

This table reports the analysis of sustainable funds' votes on environmental and social shareholder proposals. For each fund-year, we calculate the percentages of Environmental and Social (ES) proposals received by the companies held by the fund for which the fund voted "For," "Against," "Abstain," and "Do Not Vote". In Panel A, summary statistics are provided for benchmark, financial, moral, and impact funds. In Panel B, the regression results are reported for regressing the voting percentages on the financial fund dummy and impact dummy variables. *t*-stats are reported in parentheses. All regressions in Panel B include year–quarter fixed effects, with standard errors clustered at the fund level. Panel C reports the results from regressions estimated at the fund–proposal level. The dependent variable is an indicator equal to one if the fund voted Yes on the proposal. The key independent variables are Outcome Proposal, a dummy variable indicating whether the proposal is an outcome proposal; Impact Fund, a dummy variable identifying impact funds; and the interaction between Impact Fund and Outcome Proposal. A proposal is defined as an Outcome Proposal if it falls into an ISS category whose proposal descriptions (AgendaGeneralDesc and ItemDesc) are not related to ESG reporting. This definition includes proposals with the following ISSAgendaItemIDs: S0224, S0411, S0416, S0703, S0732, and S0745. Column (1) includes all sustainable funds, and column (2) includes impact funds and benchmark funds. Firm fixed effects are included, and standard errors are clustered at the proposal level.

Panel A: Summary Statistics

	(1) Benchmark	(2) Financial	(3) Moral	(4) Impact
Vote "For" ES	25.1%	40.7%	47.0%	58.5%
Vote "Against" ES	63.7%	53.1%	47.2%	34.1%
Abstain ES Vote	7.0%	2.3%	2.4%	2.8%
Do not vote	2.8%	2.1%	2.0%	1.1%
Observations	20440	4317	1164	564

Panel B: Regression Results

Panel B-1: Financial vs Benchmark

	(1) Vote "For"	(2) Vote "Against"	(3) Abstain	(4) Do not Vote
Financial Fund	11.809***	-9.587***	-2.552***	-0.341
	(8.055)	(-6.378)	(-4.540)	(-0.812)
Year-quarter FE	Yes	Yes	Yes	Yes
Clustering at Fund	Yes	Yes	Yes	Yes
Observations	24757	24757	24757	24757
Adjusted R^2	0.135	0.090	0.036	0.052

Panel B-2: Impact Fund Analysis

	(1) Vote "For"	(2) Vote "Against"	(3) Abstain	(4) Do not Vote
Impact Fund	19.809***	-20.837***	-0.147	-1.102**
	(5.265)	(-5.341)	(-0.164)	(-2.293)
Year-quarter FE	Yes	Yes	Yes	Yes
Clustering at Fund	Yes	Yes	Yes	Yes
Observations	5591	5591	5591	5591
Adjusted R^2	0.115	0.129	0.018	0.042

Panel C: Outcome Proposal Analysis

	(1) Vote "For"	(2) Vote "For"
Impact Fund*Outcome Proposal	0.181***	0.212***
	(3.326)	(3.317)
Outcome Proposal	-0.061**	-0.070***
	(-2.242)	(-3.186)
Impact Fund	0.201***	0.349***
	(19.218)	(33.393)
Sample	All Sustainable Funds	Impact and Benchmark Funds
Firm FE	Yes	Yes
Clustering at Proposal	Yes	Yes
Observations	71651	251183
Adjusted R^2	0.113	0.136

Table 8: Changes in ESG Ratings and Carbon Emission Intensity During Holding Periods

This table reports the changes in portfolio company ESG ratings and Carbon Emission Intensity during periods when firms are held by different fund types. The starting time of a stock's holding session in an impact fund is defined as the first year-month when the stock appears in the fund's holdings and the fund is classified as an impact fund at the time of entry. The holding session is defined as the consecutive months during which the stock is observed in the fund's holdings without interruption, and the fund remains classified as an impact fund throughout. The end time of the holding session is defined as the last year-month when the stock appears in the holdings. A similar definition applies to benchmark, financial, and moral funds. For funds with at least 12 (18, or 24) months of holdings, the change in ESG ratings and Carbon Emission Intensity is computed as the difference between the values 12 (18, or 24) months after the entry time and those at the entry time. Carbon intensity is interpolated between reporting years using weighted averages of last and next available data. Panel A reports the average change in ESG ratings across different fund types, while Panel B reports the average change in carbon emission intensity across different fund types.

Panel A: Changes in ESG Scores During Holding Periods

	(1) Benchmark	(2) Financial	(3) Moral	(4) Impact
12-Month ESG Score Change	0.126	0.151	0.136	0.152
12-Month E Score Change	0.165	0.143	0.111	0.153
12-Month S Score Change	0.088	0.158	0.162	0.148
Observations	1897871	258266	210784	32950

	(5) Benchmark	(6) Financial	(7) Moral	(8) Impact
18-Month ESG Score Change	0.194	0.230	0.206	0.221
18-Month E Score Change	0.246	0.222	0.187	0.260
18-Month S Score Change	0.141	0.235	0.226	0.179
Observations	1668374	234078	195567	27005

	(9) Benchmark	(10) Financial	(11) Moral	(12) Impact
24-Month ESG Score Change	0.260	0.283	0.247	0.326
24-Month E Score Change	0.320	0.278	0.200	0.343
4-Month S Score Change	0.197	0.286	0.295	0.302
Observations	1377209	201844	181546	21827

Panel B: Changes in Emission Intensity During Holding Periods

	(1) Benchmark	(2) Financial	(3) Moral	(4) Impact
12-Month Scope 1&2 Intensity Change	-4.276	-7.286	-4.148	-9.421
12-Month Scope 1&2&3 Intensity Change	-5.465	-9.251	-4.681	-13.448
Observations	2040420	254761	210948	33796

	(5) Benchmark	(6) Financial	(7) Moral	(8) Impact
18-Month Scope 1&2 Intensity Change	-6.994	-12.370	-9.347	-12.758
18-Month Scope 1&2&3 Intensity Change	-9.401	-17.221	-11.888	-16.666
Observations	1786126	227448	193925	28022

	(9) Benchmark	(10) Financial	(11) Moral	(12) Impact
24-Month Scope 1&2 Intensity Change	-10.312	-17.243	-14.765	-16.075
24-Month Scope 1&2&3 Intensity Change	-14.519	-27.100	-21.820	-20.241
Observations	1447923	193831	178498	22452

Figure 1: Sources of Asset Growth in Financial, Moral, and Impact Funds, 2014–2023

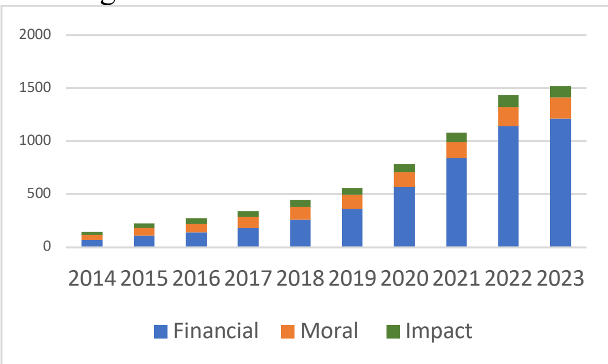

Panel A-1: Total Fund Count

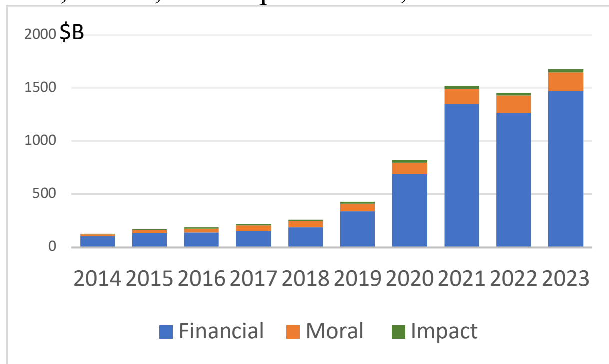

Panel A-2: Total Fund AUM ($B)

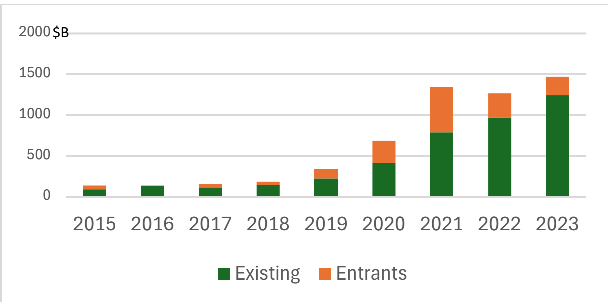

Panel B-1: Financial Funds—AUM and Decomposition ($B)

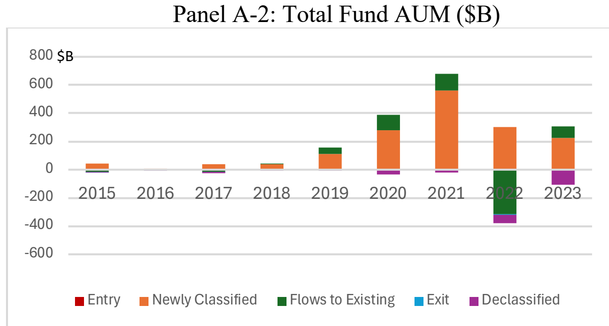

Panel C-1: Financial Funds—ΔAUM and Decomposition ($B)

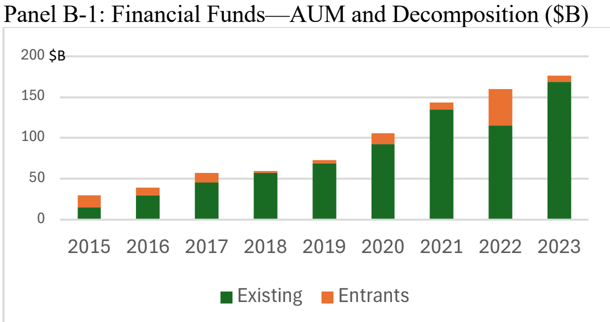

Panel B-2: Moral Funds—AUM and Decomposition ($B)

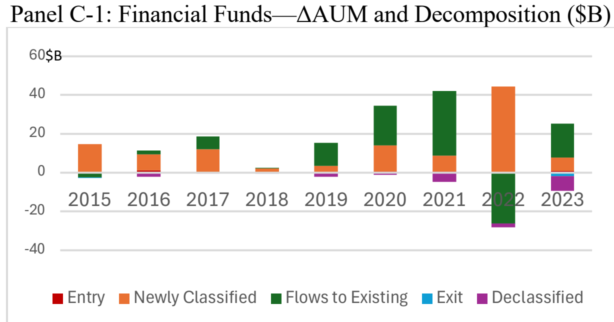

Panel C-2: Moral Funds—ΔAUM and Decomposition ($B)

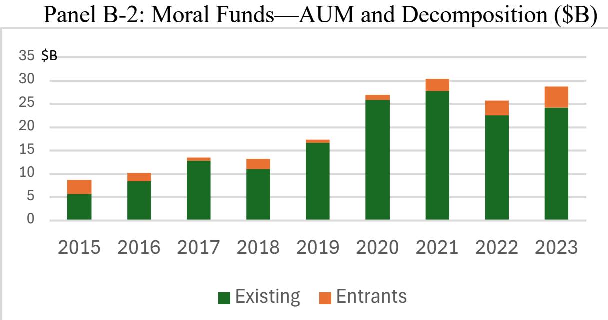

Panel B-3: Impact Funds—AUM and Decomposition ($B)

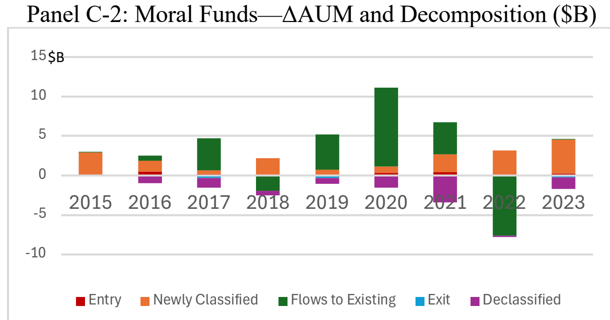

Panel C-3: Impact Funds—ΔAUM and Decomposition ($B)

This figure plots the level, composition, and year-over-year change in assets under management (AUM) for financial, moral, and impact funds from 2014 to 2023. Panel A-1 reports the number of funds by category in each year, and Panel A-2 reports the total AUM by category. Panels B-1 to B-3 report total AUM within each category (financial, moral, and impact) and decompose it into AUM of funds classified in prior years ("Existing") and AUM of funds that are either newly incepted or newly classified in the current year ("Entrants"). Panels C-1 to C-3 report the year-over-year change in AUM within each category and decompose it into contributions from newly created funds ("Entry"), newly classified funds ("Newly Classified"), changes in AUM of previously classified funds ("Flows to Existing"), and decreases due to fund exit ("Exit") or declassification ("Declassified").

Figure 2: Changes in Prospectus Sustainability Language around Classification Events

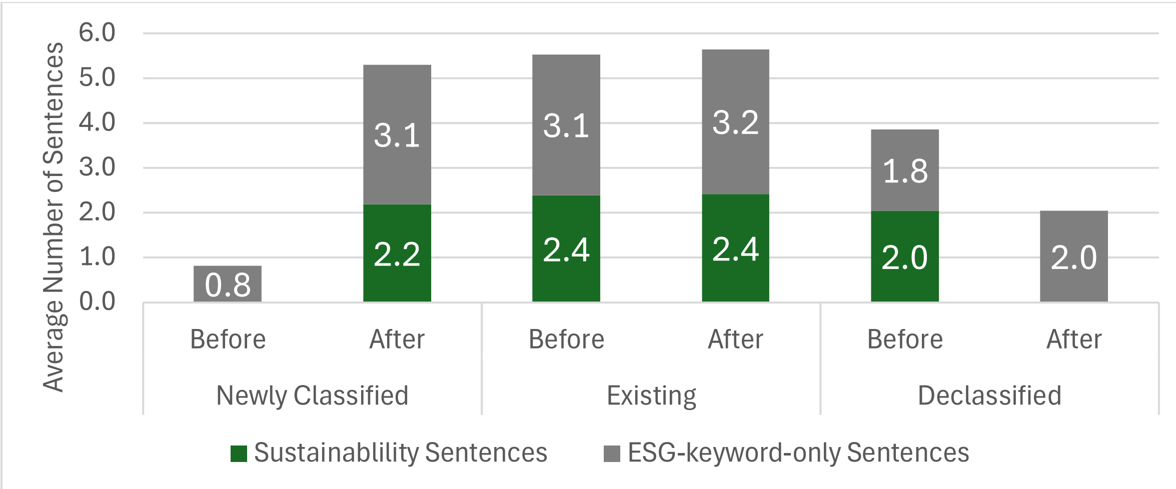

Panel A: Sustainability Language Around Classification Changes

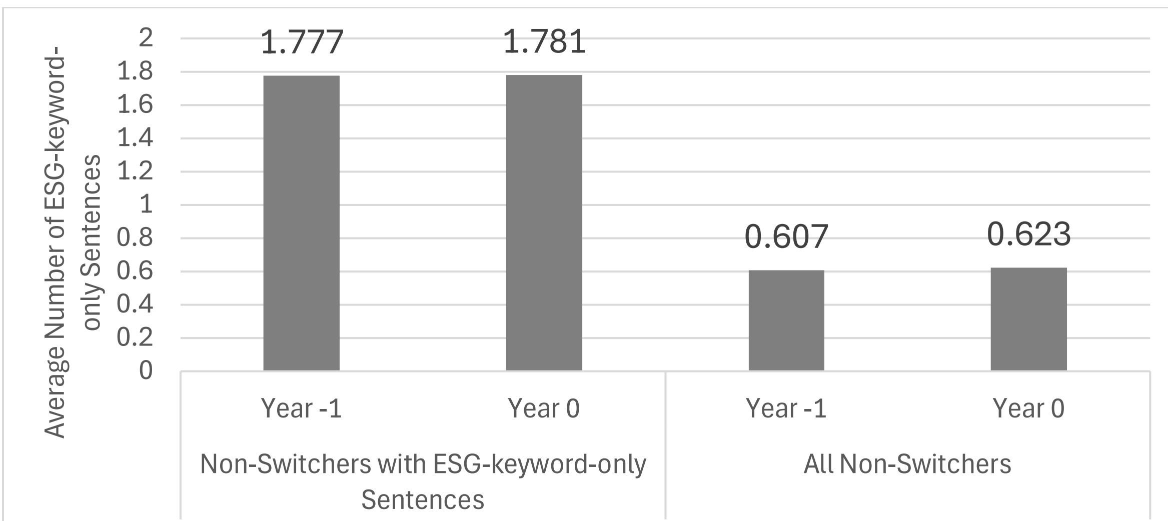

Panel B: Counterfactual Prospectus Language for Non-Sustainable Funds

This figure plots the average number of Sustainability Sentences and ESG-Keywords-Only Sentences across fund groups. Sustainability Sentences are prospectus sentences classified as *financial*, *moral*, or *impact*. ESG-Keywords-Only Sentences contain ESG keywords but are not classified as Sustainability Sentences. Panel A reports average sentence counts for funds around changes in sustainability classification status. For newly classified sustainable funds, we report averages in the year prior to classification (“Before”) and the year of classification (“After”). We also report averages for funds that remain sustainable in both periods and for funds that are declassified in the current year. Panel B reports weighted averages for non-sustainable funds. For each year, we compute average sentence counts in the current and prior year separately for non-sustainable funds with and without ESG-Keywords-Only sentences and weight these averages by the number of newly classified sustainable funds in that year. This weighting provides a counterfactual benchmark, capturing how prospectus language evolves over time for non-switching funds absent a change in sustainability classification.

Figure 3: Sector Weights, ESG Ratings, and Emission Intensity

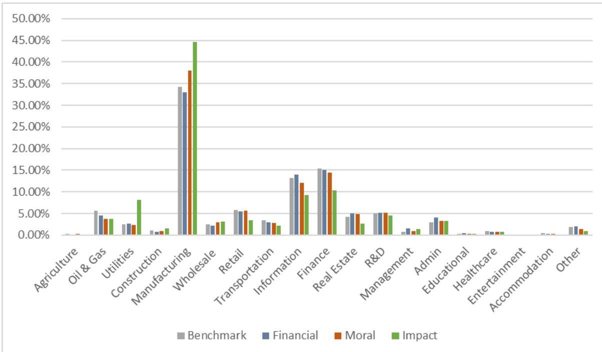

Panel A: Sector Weights

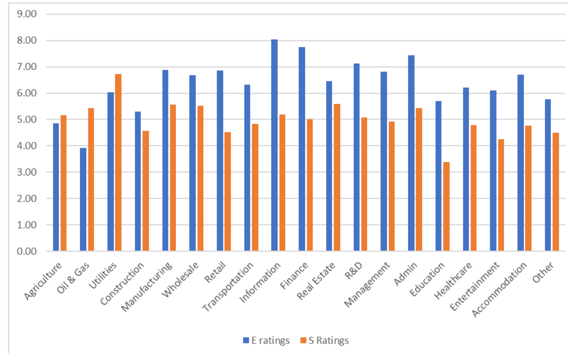

Panel B: Sector Average E and S Ratings

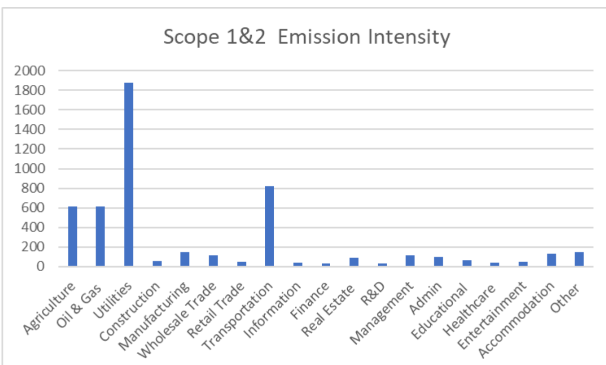

Panel C: Scope 1&2 Emission Intensity

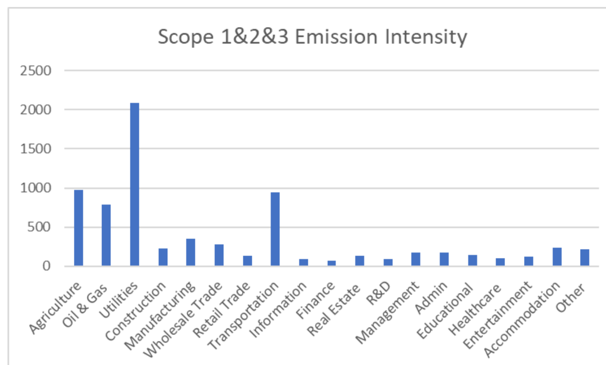

Panel D: Scope 1&2&3 Emission Intensity

This figure compares portfolio sector weights of sustainable fund holdings by type and sector-average E and S ratings and emission intensity of covered firms by MSCI and Trucost, respectively. Panel A shows the sector portfolio weights for benchmark, financial, moral, and impact funds. Panel B shows the average MSCI covered firms' E and S ratings by sector. Panel C and D show the average Trucost covered firms' scope 1&2 and 1&2&3 emission intensity by sector, respectively.

Figure 4: DiD Analysis of Firm Emission Intensity during Holding Periods

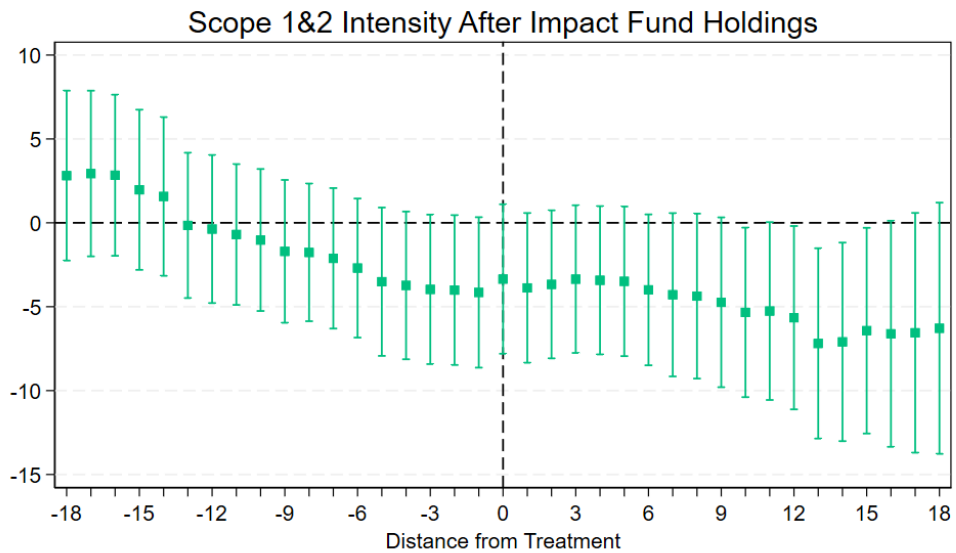

Panel A: Impact Fund Holdings

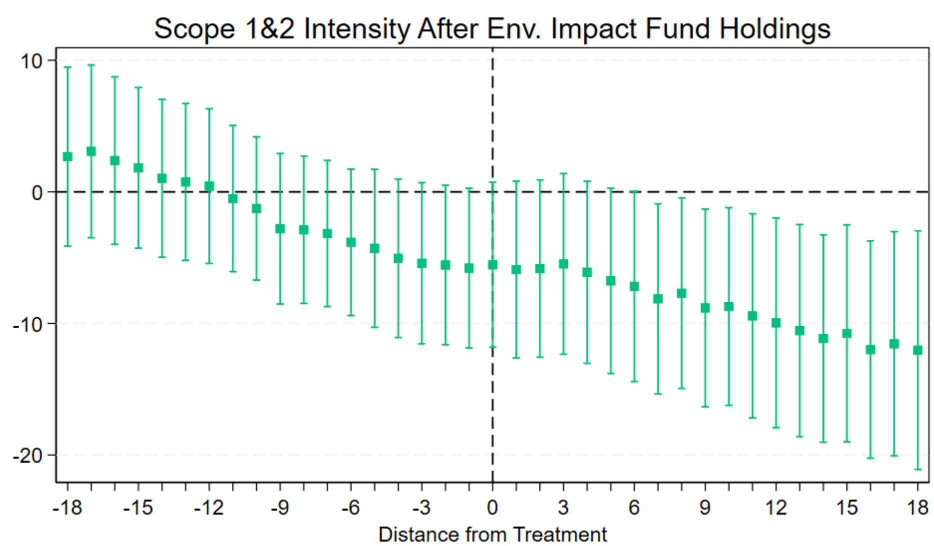

Panel B: Environmental Impact Fund Holdings

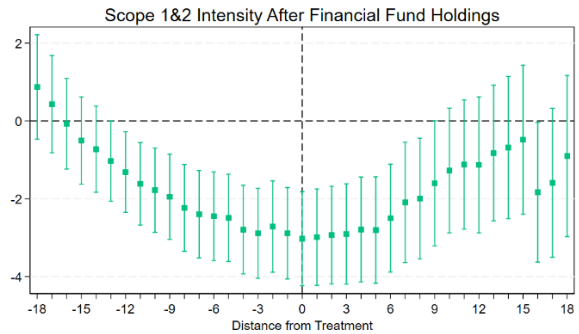

Panel C: Financial Fund Holdings

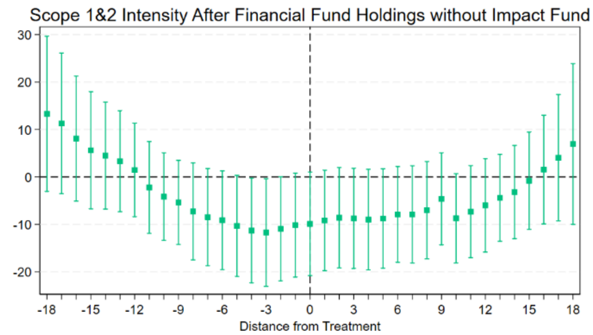

Panel D: Financial Fund Holdings with No Overlapping Impact Fund Holdings

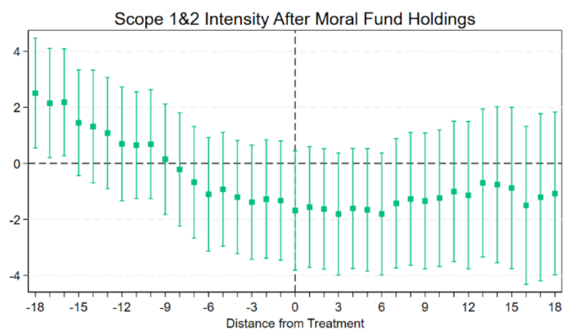

Panel E: Moral Fund Holdings

This figure shows the dynamic treatment effect on firms' Scope 1 & 2 carbon intensity after acquisition by sustainable funds. We estimate a dynamic treatment effect model.

$$y_{i,t} = \sum_{k=-18}^{18} Holding_{i,t}^{k} + FEs + \varepsilon_{i,t}$$

Where $Holding_{i,t}^{k}$ is a dummy variable equal to one for firm i at calendar time t if t is k months before/after the first month in which firm i becomes held by a fund of a given sustainability type for $k \in [-18,18]$. For each treated firm—defined as a stock newly purchased by a sustainable fund—we identify a matched control firm based on ROA and log(total assets) within the same 2-digit NAICS sector. Control firms are further required to have received no sustainable fund investment between t–24 and t+24. The outcome variable is Scope 1&2 carbon intensity. Firm and year-month fixed effects are included, and standard errors are clustered at the firm level. Confidence intervals are reported at the 90% level. Panel A reports results for impact fund holdings, Panel B reports environmental impact fund holdings, focusing on funds in the top quintile based on their support for environmental related shareholder proposals, Panel C reports results for financial fund holdings, Panel D reports results for financial fund holdings without firms that are also held by impact funds between *t*-18 to *t*+18, and Panel E reports results for moral fund holdings.

Internet Appendix for

Decoding Sustainable Investment Strategies: Bridging Intentions and Outcomes

Table of Contents

IA1. ESG Sentence Classification

IA1.1 Financial Sentences

The first category is Financial. A sentence is classified as Financial if it indicates that the fund employs ESG information to enhance financial performance. In general, sentences in this category share the following characteristics.

(1) direct references to using using ESG for financial performance

"...but the firm believes that responsible corporate behavior with respect to ESG factors can contribute to positive and sustainable long-term financial performance."

(2) direct references to using using ESG as risk factor

"...evaluating ESG-related risks as part of its research recommendations."

(3) direct references to using using ESG as material issues

"The fund will not invest in companies that Newton deems to have material environmental, social or governance issues.."

(4) key terms indicating ESG use for financial performance, not externalities

- ESG as attractive investment attributes
- ESG as Sustainable business practices
- ESG controversies

  "Newton may also monitor certain companies in whose securities the fund has invested for emerging environmental, social or governance controversies and issues and may update a company's ESG quality review rating on the basis of such monitoring."

- ESG ratings

  "As part of this research process, Invesco may use third-party ESG ratings, company reporting, and engagement with management."

- ESG issues

  "The Fund's Adviser may consider information about environmental, social and governance issues (also referred to as ESG) in its bottom-up stock selection process when making investment decisions."

(5) Index indicating Financial motives

- ESG MSCI

- ESG Sustainalytics US Inc.

(6) use ESG information as an input when making investment selections

"As such, the Adviser evaluates ESG factors as part of the investment analysis process and this forms an integral component of the Adviser's quality rating for all companies."

Investment analysis process suggests that ESG information is used for financial performance.

(7) references to companies benefiting from ESG transitions

"Stocks are included in the Underlying Index based on the Index Provider's evaluation that such companies will substantially benefit from a societal transition toward the use of cleaner energy, zero-CO_2 renewables, and conservation."

IA1.2 Moral Sentences

The second category is Moral. A sentence is classified as Moral if it indicates that the fund excludes certain types of companies. We assume that such exclusions are based on ethical rather than financial considerations. A sentence is not labeled as Moral if it explicitly states that exclusions are made for non-Moral reasons. In our manual review, we treat exclusion criteria as Moral-related unless they are specifically described as non-Moral.

"The Fund is fossil-fuel free, as it does not invest in companies that derive significant revenues from the extraction, exploration, production or refining of fossil fuels; the Fund may invest in companies that use fossil fuel-based energy to power their operations or for other purposes."

IA1.3 Impact Sentences

The third category is Impact. A sentence is classified as Impact if it indicates that the fund uses ESG information specifically for externality considerations, without reference to financial performance. Specifically referencing externality considerations without mentioning financial performance, and indicating that the fund values impact in its own right.

(1) specifically referencing externality considerations without mentioning financial performance, and indicating that the fund values impact in its own right.

"Rockefeller's sustainability and impact evaluation considers environmental, social, and governance criteria such as corporate governance practices, product quality and safety, workplace diversity practices, environmental impact and sustainability, community investment and development, and human rights record."

"Strictly in accordance with its guidelines and mandated procedures, WilderShares, LLC (the "Index Provider") compiles and maintains the Underlying Index, which is composed of stocks of publicly traded companies listed on a major exchange in the United States that are engaged in the business of the advancement of cleaner energy and conservation or are important to the development of clean energy."

(2) directly stating the creation of positive impact or the measurement of impact, without providing specific details.

“ The Fund may invest a significant amount of its assets in taxable and tax-exempt municipal bonds that finance community projects whose primary purpose, in the Adviser’s view, is a positive impact to the community in which the project is located.”

“The Adviser researches the use of proceeds of each bond invested in by the Fund utilizing both quantitative metrics and qualitative details to measure the impact achieved.”

IA1.4 Unclassified Sentences

The Unclassified category includes sentences that are too ambiguous to be clearly assigned to a single category, as well as sentences containing generic ESG content that does not reflect any specific investor philosophy. A few typical examples are provided below.

“The Fund seeks to achieve its investment objective while applying an environmental, social and governance (ESG methodology developed by the Fund’s subadviser in the selection of portfolio investments.)”

“ESG considerations are fully integrated across all asset classes.”

“Final investment decisions are made by the portfolio managers typically on the basis of market conditions as well as technical and ESG considerations with respect to both specific instruments and the overall composition of the portfolio.”

IA1.5 Fund-Level Classification and Tie-Breaking Rule

In fewer than 2% of fund-quarter observations, two (or three) categories are tied for the largest sentence share. In such cases, the fund is assigned to each tied category. For example, a fund tied as financial and moral is included in both comparisons against benchmarks and carries both dummies in regressions comparing financial and moral funds.

IA2. Manually Classified Sample Construction

In this section, we provide a detailed explanation of the process used to construct the manually classified training and testing sample. We begin with a list of U.S. sustainable mutual funds compiled and published by Morningstar. Morningstar began compiling this annual list in 2018, and we obtained the lists for 2018, 2019, 2020, and 2022 either directly from the company’s website or from other researchers who have used these data. We focus on this set of U.S. mutual funds identified by Morningstar as pursuing sustainability goals because we expect these funds to be more likely than other funds to describe their sustainability objectives in the investment strategy section of their prospectuses. Oversampling true positives, defined as sentences classified as Financial, Moral, or Impact, from imbalanced samples is a common strategy in machine learning and textual analysis to improve classification balance and performance (He and Garcia, 2009). Using a random number generator, we select 2,834 sentences for manual classification.

To ensure that our manually classified sentences also include potential sustainable funds outside the Morningstar list, we supplement the sample with 741 sentences drawn from funds not listed by Morningstar. This additional sample from Non-Morningstar funds also oversamples true positives and follows standard practices in machine learning and textual analysis.

In total, our manually classified sample is drawn from 362 fund prospectuses and includes 3,575 ESG-related sentences. We then randomly divide the sample into 2,950 sentences for training and 625 sentences for testing.

IA3. Comparison with Generative AI

IA3.1 Model Comparison: LDA, BERT, GenAI

In this section, we compare the predictions of our BERT model with those of a generative AI (GenAI) model. One natural methodological question, given that GenAI models are widely used for daily tasks, is why we do not rely on a GenAI model for classification. To address this question, we first briefly discuss the methodological differences among three potential textual analysis approaches, with a primary focus on comparing the BERT model and the generative AI model. Table IA1 provides a summary of the key differences across model types.

[Insert Table IA1: Language Model Comparison]

One possible classification approach is topic modeling, such as Latent Dirichlet Allocation (LDA). These methods are conceptually similar to principal component analysis in that they allow the data itself to be represented by a set of latent topics (i.e., they are unsupervised), with each topic potentially corresponding to a class. However, models of this type are limited in their ability to capture complex relationships between words. Moreover, they are not trained on ground-truth labels, meaning we cannot directly assess whether the learned topics correspond to correct human coding or generate classifications that align with the true labels. Topic models are not well suited to our setting because the subtlety of sustainable investment motives requires a model that can understand and process complex relationships between words and can be trained to learn the ground-truth classifications identified by researchers.

The BERT model is the approach used in our paper. It has the advantage of being able to capture complex relationships between words through its transformer architecture. BERT is pre-trained so that the base model can already infer word meanings from their full context. Moreover, it is a supervised model and can be trained on ground-truth labels provided by researchers. By fine-tuning the model using researcher-labeled data, BERT can generate classifications that align closely with the ground truth, combining the benefits of rich contextual language representations from pre-training with task-specific learning from supervised re-training.

One alternative classification approach is to use a Generative AI (GenAI) model. Modern GenAI models—such as ChatGPT, Google's Gemini series, and Claude—are built on transformer architectures. While these models are substantially larger than BERT, they share a similar architectural foundation and are capable of modeling complex relationships between words and phrases. Because these models are pretrained on massive text corpora, one could directly prompt them to classify sustainable investment goals into financial, moral, and impact categories based on the model's semantic understanding of these concepts. In this setting, however, the training data underlying the model implicitly serves as the "ground truth," rather than labels explicitly defined and provided by the researcher. For example, when ChatGPT is asked to classify funds based on its reading of fund prospectuses, the classification is generated without supervision and without being trained or calibrated on researcher-provided ground-truth labels. This methodology is problematic because retail investors and the media have been documented to often misunderstand sustainable goals, and such misunderstandings reflected in texts appearing in news articles and online forum discussions may be transmitted to the model through its pretraining data. In the next subsection, we present the results obtained using this approach and compare them with our BERT-based classification.

One way to "force" a GenAI model to produce outputs closer to researcher-defined ground truth is through carefully designed prompting or through Retrieval-Augmented Generation (RAG). Prompt engineering can steer the model's responses toward the categories intended by the researcher; for example, when interacting with ChatGPT, prompts are often structured to constrain the model to answer a specific question in a desired format. RAG further extends this idea by augmenting prompting with an explicit retrieval step, in which researcher-provided ground-truth examples are supplied to the model at inference time, thereby encouraging classifications that align more closely with the reference labels.

However, this approach is not an ideal implementation of GenAI for classification tasks. If the objective is to ground the model in human researchers' labeled training samples, a supervised model such as BERT can directly learn this mapping through fine-tuning, achieving satisfactory performance with substantially greater stability and without hallucinations. Consequently, when researcher-defined ground truth is available for training, GenAI models become a less suitable choice relative to BERT-based classifiers.

IA3.2 GPT 5 Results

In this section, we compare the classifications generated by the BERT model with those produced by a generative AI model. Specifically, we use OpenAI's GPT-5 model for comparison. We adopt "gpt-5-2025-08-07" as a representative snapshot to benchmark BERT model performance against that of a ChatGPT-style model. The "gpt-5-2025-08-07" model is a fixed snapshot of GPT-5 released on August 7, 2025. It is designed for general-purpose coding and chatbot tasks and is widely used and regarded as highly successful. We use the following prompt for classification:

"You are an investor reading a mutual fund prospectus. Given a sentence related to ESG (Environmental, Social, and Governance), classify the fund's investment approach based on its primary intent: Financial – primarily focused on generating financial returns; Moral – guided by ethical or moral values; Impact – aiming to achieve measurable environmental or social outcomes; Unclear – if there isn't enough information to make a clear classification. Respond with only one word: 'Financial', 'Moral', 'Impact', or 'Unclear'."

[Insert Table IA2: ChatGPT Model Performance]

Table IA2 provides the ChatGPT model prediction performance. We report the performance of ChatGPT-5 prediction performance on the 625 sentences in the testing sample. We find that ChatGPT-5 prediction performance is significantly lower than BERT model. The accuracy for financial sentences is 0.72 with ChatGPT-5, compared to 0.91 in BERT model, which is a 19-percentage points decreases. The Moral and Impact accuracy are also smaller. We find that the decrease in performance is mainly due to the low precision of ChatGPT-5 prediction, in Financial, Moral, and Impact prediction, most strongly for financial sentences. The Precision decreases from 0.85 to 0.47, from 0.85 to 0.62, from 0.83 to 0.61, for Financial, Moral, and Impact Respective. Too many false positives are given with ChatGPT prediction. Overall, we find that ChatGPT-5 underperforms our BERT model.

IA4. Textual Analysis Background

BERT is able to capture relationships between words because it is trained to model contextual dependencies through self-attention. BERT uses a Transformer architecture in which each word attends to every other word in the sentence. This self-attention mechanism allows the model to weight the relevance of surrounding words when forming a representation for a given token, enabling it to capture semantic, syntactic, and long-range relationships. During pretraining, BERT is optimized using a masked language modeling objective, which forces the model to infer a missing word from its full context. As a result, the learned representations encode rich information about how words relate to each other across different contexts, allowing BERT to distinguish meaning based on usage rather than relying solely on individual word frequencies.

BERT fine-tuning and retrieval-augmented generation (RAG) represent fundamentally different approaches to textual analysis. Fine-tuning a BERT model involves supervised learning in which model parameters are updated using labeled data to directly optimize performance on a specific task, such as text classification. As a result, task-relevant information is embedded in the model's parameters, yielding stable, reproducible predictions that are well suited for empirical analysis and hypothesis testing. In contrast, RAG-based large language models do not learn task-specific decision rules from labeled data. Instead, they retrieve relevant external documents at inference time and generate free-form responses conditioned on the retrieved text and the prompt. While RAG systems are flexible and effective for open-ended question answering, their outputs are inherently prompt-dependent and stochastic, and they lack a fixed decision boundary. Consequently, RAG is not a substitute for fine-tuned discriminative models when the research objective requires consistent classification, interpretability, and reproducibility.

Prompting and retrieval-augmented generation (RAG) are related but conceptually distinct. A prompt refers to the textual input used to condition a large language model's behavior at inference time and relies entirely on the model's pretrained, parametric knowledge. In contrast, RAG is a system-level architecture that augments prompting with an explicit retrieval step, in which relevant external documents are retrieved and incorporated into the prompt before generation. While both approaches operate without updating model parameters, RAG enables the model to condition its output on non-parametric, externally stored information, whereas prompting alone does not involve retrieval or access to external knowledge. Consequently, prompting is a component of AG, but RAG represents a broader framework designed to extend language models beyond their fixed internal knowledge.

IA5. Discussion of Additional Results

IA5.1 Sector Exclusion by Moral Funds

We examine whether moral funds that express exclusionary investment goals in their prospectuses indeed underweight sectors they claim to avoid. In the manual coding of prospectus sentences, we find that alignment of investment goals with moral or ethical values is frequently expressed as prohibition against certain business activities and exclusion of companies engaged significantly in such activities. Some of the excluded activities reflect traditional religious values (e.g., alcohol, gambling, weapons, and abortion), while others are grounded in modern ethical or scientific concerns (e.g., fossil fuels and climate impact, tobacco and health risks). Moral funds vary in orientation—some are religious (e.g., Catholic, Presbyterian), while others are secular—leading to heterogeneity in exclusion practices. For example, some funds may avoid casinos but not coal, while others exclude fossil fuels and weapons but not abortion providers. As a result, average portfolio weights in excluded sectors are expected to be lower than those of benchmark funds, but not zero.

We use the following NAICS sectors corresponding to commonly excluded categories for this analysis: "Tobacco", "Oil & Gas Extraction", "Natural Gas Distribution", "Coal Mining", "Casino & Gambling", and "Aerospace" (including weapons manufacturers). These variables are defined in Appendix A2.

Panel A of Table IA3 presents the sector weights of fund portfolios in these categories by benchmark, sustainable, and moral funds. As expected, moral funds have substantially lower portfolio weights in these sectors than benchmark funds.

[Insert Table IA3: Sector Exclusion by Moral Funds]

In Panel B, we report the results of regression analysis. We find that moral funds hold significantly smaller portion of their portfolios in tobacco, gas, casino and aerospace stocks relative to benchmark funds in the same quarter-year. The coefficients for oil & gas and coal mining are also negative but insignificant, possibly because traditional religious funds—which make up a substantial share of our sample—typically do not exclude fossil fuels. We plan to collect fund-specific exclusion categories to enable more granular analysis and sharper inferences.

IA5.2 Environmental Rating Decomposition

To further examine the heterogeneity of sustainable funds' investment strategies, we decompose the overall environmental (E) ratings into four underlying MSCI subcategories and repeat the comparison analysis from Table IA4.

[Insert Table IA4: Environmental Rating Decomposition]

Panel A of Table IA4 reports the value-weighted average scores for the four MSCI environmental subcategories—**Climate Change**, **Natural Capital**, **Pollution & Waste**, and **Environmental Opportunities**—for benchmark, sustainable, financial, and impact funds. Among benchmark funds, the highest scores are observed in the Climate Change category, while the lowest appear in Environmental Opportunities.

According to MSCI's ESG Ratings Methodology Manual, the *Climate Change* subcategory captures companies' exposure to operational and supply chain risks stemming from carbon pricing, regulatory emissions limits, and physical climate hazards. *Natural Capital* assesses exposure to biodiversity loss, deforestation, unsustainable raw material sourcing (e.g., palm oil), and water stress. *Pollution & Waste* reflects potential liabilities from regulatory actions related to e-waste, packaging, toxic emissions, and contamination. Finally, *Environmental Opportunities* evaluates how well-positioned firms are to benefit from the transition to a low-carbon economy—particularly in clean technology, green buildings, and renewable energy.

Comparing fund types, we find that financial funds generally hold stocks with higher environmental subcategory scores than benchmark funds across all four dimensions except *Pollution & Waste*. In contrast, impact funds hold stocks with higher *Pollution & Waste* and *Environmental Opportunities* scores relative to financial funds.

Panel B presents regression results controlling for time fixed effects. In Panel B-1, financial funds are shown to hold firms with significantly higher scores in all four subcategories compared to benchmark funds, consistent with their preference for ESG "leaders" as a means of managing material risk and enhancing returns. Panel B-2 compares impact funds to financial funds. We find that impact funds hold stocks with significantly higher *Pollution & Waste* and *Environmental Opportunities* scores—consistent with investing in companies whose core businesses (e.g., renewable energy, green buildings, and clean tech) contribute directly to environmental solutions. These represent a class of investments where the business model itself serves as a generator of positive environmental impact.

However, we also find that impact funds hold companies with significantly lower *Natural Capital* scores relative to financial funds. These firms are more likely to be associated with operations that generate negative externalities—such as deforestation, water depletion, or unsustainable sourcing—and thus face regulatory, physical, or reputational risks. While MSCI ratings may only imperfectly capture the full extent of these externalities, this finding is consistent with the interpretation that impact funds allocate capital to relatively poor environmental performers, potentially as part of an engagement-based strategy aimed at improvement. Notably, the higher *Environmental Opportunities* scores among impact fund holdings are also consistent with our earlier finding that these funds tilt more heavily toward the utilities sector (see Table 2 and Figure 3), where opportunities for energy transition are more prevalent.

Taken together, the decomposition analysis reinforces the view that impact funds differ fundamentally from financial funds in how they approach sustainability. Impact funds allocate capital to firms that are both positioned to reduce environmental harms (e.g., by improving supply chains or reducing emissions) and to deliver environmental solutions through their core products and services. In contrast, financial funds predominantly allocate to firms that already score well on environmental metrics across the board. These differences reflect divergent approaches: one emphasizing risk-adjusted return enhancement through ESG integration, the other emphasizing real-world impact generation through improvement and innovation.

IA5.3 Robust Flow Performance Results

This section examines whether the results on the association between mutual funds' stated sustainable investment strategies and flow–performance sensitivity are robust to the use of alternative decay functions. We show that the results remain robust when using different decay functions and alternative fixed effects.

Table IA5 investigates how flows respond to performance. Following Barber, Huang, and Odean (2016), we construct alphas using an exponential decay function applied to prior monthly returns, with the CAPM alpha serving as the performance metric given its stronger predictive power for flows. Panel A compares moral and benchmark funds. Across specifications, we find that moral funds earn significantly higher unconditional flows (roughly 0.3% per month). More notably, moral funds exhibit significantly attenuated flow–performance sensitivity. For example, in column (2), a 1 percentage point decline in alpha reduces flows by 0.388% for benchmark funds, but by only 0.195% (= 0.388 − 0.193) for moral funds. Using the 18-month decay in column (4), the analogous decline is 0.823% for benchmark funds versus 0.499% for moral funds. These findings are consistent with moral funds attracting investors who derive utility from ethical exclusionary screens and who are less responsive to performance.

[Insert Table IA5: Flow Performance Sensitivity]

Panel B compares impact and benchmark funds. Unlike moral funds, impact funds do not exhibit statistically different flow–performance sensitivity relative to benchmarks. Although impact funds display significantly higher unconditional flows in columns (1) and (2), the difference becomes insignificant when using the 18-month alpha decay (columns (3)–(4)). Overall, we find no evidence that impact funds are systematically matched with less performance-sensitive investors.

Panel C compares financial and benchmark funds. As expected, we find no statistically significant difference in flow–performance sensitivity between the two groups. This is consistent with financial funds competing for flows from traditional pecuniary investors.

Our findings underscore the importance of distinguishing among sustainable funds by stated objectives. Our results are robust when using an alternative three-month decay function and when excluding fixed effects.

IA5.4 Voting Behavioral of Moral Funds

To further unpack the heterogeneity among moral funds, we classify them into subgroups based on the presence of hybrid objectives: *pure moral*, *financial-moral*, *impact-moral*, and *financial-impact-moral*.[1] Panel A-2 presents the average support rates for each subgroup. We find that *pure moral* funds—those with no hybrid classification—exhibit voting behavior nearly identical to financial funds. In contrast, moral funds with impact-related language (e.g., *impact-moral*, *financial-impact-moral*) support ES proposals at substantially higher rates. These findings suggest that moral funds with impact elements adopt more activist voting behavior, while pure moral funds behave similarly to financially motivated funds in the proxy voting context.

[Insert Table IA6: Moral Fund Votes on Environmental and Social Shareholder Proposals]

[1] As described in the methodology section, funds are classified into one of three types—financial, moral, or impact—based on the most frequent classification among their sustainability-related sentences. For example, a fund is classified as "moral" if the number of sentences labeled as moral exceeds those labeled as financial or impact. Within the set of moral funds, we further define subgroups such as *financial-moral* (moral funds with at least one sentence also classified as financial) and *impact-moral* (moral funds with at least one sentence also classified as impact), and so on.

IA5.5 Emission Intensity Changes: Difference-in-Differences Analysis

We report the coefficients used to construct Figure 4, which presents the difference-in-differences analysis. We also provide results for Scope 1&2&3 emission intensity. The results are reported in Table IA7. We find very similar results, with no evidence of pretrends, and treated firms exhibiting significantly lower emission intensity after treatment.

[Insert Table IA7: Difference-in-Differences Analysis of Emission Intensity]

In our main specification, if two funds begin holding firm i at the same time, we still set the indicator equal to one. In Figure IA3, we present robustness results in which, in such cases, we instead replace the indicator with the total number of simultaneous holding events that share the same entry date. We find very similar results, with no evidence of pretrends and treated firms exhibiting significantly lower emissions intensity after treatment.

[Insert Figure IA3: Robust DiD Analysis of Firm Emission Intensity during Holding Periods]

IA5.6 Shareholder Proposal Measures

This section describes our methodology for constructing environmental and social (ES) shareholder proposal measures. We identify ES proposals using ISS category codes based on Table A1 in He, Kahraman, and Lowry (2023), supplemented by manual review. To construct outcome-focused proposals, we examine the resolution type and agenda general description for each ISS category code and review ISS's detailed descriptions. Because many ISS categories include both disclosure-oriented and outcome-oriented components, we identify six ISS categories that are exclusively outcome-oriented and use these as our measure of pure outcome-oriented proposals.

Table IA1: Language Model Comparison

Model Type	Topics (LDA) model	BERT model	Generative AI
Process complex relationship between words	No	Yes	Yes
Supervised	No	Yes	No (but prompt and RAG)
Trained on “Ground truth”	No	Yes	No (but prompt and RAG)
Model output stable over time	N/A	Yes	Yes/No
Prone to deviation and hallucination from “ground truth”	N/A	No	Yes

Table IA2: ChatGPT Model Performance

This table reports the performance of the BERT model trained to classify ESG-related sentences. We report sentence-level prediction performance using the "gpt-5-2025-08-07" model with the following prompt: "You are an investor reading a mutual fund prospectus. Given a sentence related to ESG (Environmental, Social, and Governance), classify the fund's investment approach based on its primary intent: Financial—primarily focused on generating financial returns; Moral—guided by ethical or moral values; Impact—aiming to achieve measurable environmental or social outcomes; Unclear—if there is insufficient information to make a clear classification. Respond with only one word: 'Financial,' 'Moral,' 'Impact,' or 'Unclear.'" We compare the ChatGPT-generated classifications with manual classifications from a testing sample of 625 sentences. Four different model performance measures are calculated to measure the accuracy of BERT classification. Accuracy is the ratio of (true positives + true negatives) divided by the total number of observations (fraction of correct classifications). Precision is the ratio of true positives divided by the sum of true positives and false positives. Recall is the ratio of true positives divided by the sum of true positives and false negatives. *f1* is defined as $[\frac{2}{\frac{1}{Precision}+\frac{1}{Recall}}]$. We also provide the results from BERT model for comparison.

	ChatGPT-5				BERT			
	Accuracy	Precision	Recall	f1	Accuracy	Precision	Recall	f1
Financial	0.7216	0.4749	0.764	0.5857	0.9088	0.8467	0.7888	0.8167
Moral	0.96	0.6182	0.8947	0.7312	0.9840	0.8500	0.8947	0.8718
Impact	0.912	0.6125	0.6712	0.6405	0.9440	0.8276	0.6575	0.7328

Table IA3: Sector Exclusion by Moral Funds

This table reports the portfolio weights of moral funds in sectors commonly excluded on ethical grounds. In Panel A presents mean sector weights for benchmark, sustainable, and moral funds. In Panel B, the regression results are reported for regressing the fund sector weights on the moral dummy variables for a combined sample of benchmark funds and moral funds with year-quarter fixed effects and standard errors clustered at the fund level. *t*-stats are reported in parentheses.

Panel A: Excluded Sector Weights

	Benchmark	Sustainable	Moral
Tobacco	0.38%	0.13%	0.09%
Oil & Gas Extraction	1.11%	0.96%	0.85%
Gas Distribution	0.35%	0.42%	0.25%
Coal Mining	0.03%	0.02%	0.03%
Casino & Gambling	0.05%	0.03%	0.01%
Aerospace	0.61%	0.40%	0.31%
Observations	59423	16463	3492

Panel B: Regression Analysis

	(1) Tobacco	(2) Oil & Gas	(3) Gas Distribution	(4) Coal Mining	(5) Casino & Gambling	(6) Aerospace
Moral Fund	-0.247***	-0.162	-0.094***	-0.013	-0.028***	-0.315***
	(-6.420)	(-1.572)	(-2.762)	(-0.969)	(-8.477)	(-4.763)
Observations	62915	62915	62915	62915	62915	62915
Adjusted R^2	0.013	0.042	0.004	0.005	0.022	0.007
Year-quarter FE	Yes	Yes	Yes	Yes	Yes	Yes

Table IA4: Environmental Rating Decomposition

This table reports the decomposition of the environmental (E) ratings of portfolio holdings by fund type. Subcategory ratings are from MSCI and cover Climate Change, Natural Capital, Pollution & Waste, and Environmental Opportunities. For each fund-quarter, we calculate the value-weighted Subcategory E ratings for all MSCI-rated stock holdings. In Panel A, the mean Climate Change, Natural Capital, Pollution & Waste, and Environmental Opportunities ratings are provided for the benchmark funds, sustainable funds, financial funds, and impact funds. In Panel B, the regression results are reported for regressing the fund-quarter Subcategory E ratings on the financial dummy and impact dummy variables. *t*-stats are reported in parentheses. All regressions include year–quarter fixed effects, with standard errors clustered at the fund level where indicated.

Panel A: Average Environmental Ratings

	(1)	(2)	(3)	(4)
	Benchmark	Sustainable	Financial	Impact
Climate Change	7.251	7.750	7.804	7.676
Natural Capital	6.272	6.611	6.692	6.323
Pollution & Waste	5.333	5.094	5.019	5.505
Environmental Opportunities	4.479	4.729	4.678	5.299
Observations	59413	16459	12245	1789

Panel B: Regression Results

Panel B-1: Financial Funds vs Benchmark Funds

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Climate Change		Natural Capital		Pollution & Waste		Env. Opportunities	
Financial	0.209***	0.209***	0.199***	0.199***	0.278***	0.278***	0.104***	0.104***
Fund	(15.850)	(4.148)	(13.588)	(3.484)	(20.654)	(5.949)	(12.331)	(3.172)
Year-quarter FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Clustering at Fund	No	Yes	No	Yes	No	Yes	No	Yes
Observations	71510	71510	71145	71145	67715	67715	67321	67321
Adjusted R^2	0.137	0.137	0.068	0.068	0.377	0.377	0.042	0.042

Panel B-2: Impact Funds vs Financial Funds

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Climate Change		Natural Capital		Pollution & Waste		Env. Opportunities	
Impact	0.029	0.029	-0.297***	-0.297**	0.214***	0.214**	0.678***	0.678***
Fund	(0.996)	(0.243)	(-8.700)	(-2.295)	(6.672)	(2.009)	(27.508)	(6.018)
Year-quarter FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Clustering at Fund	No	Yes	No	Yes	No	Yes	No	Yes
Observations	13789	13789	13706	13706	12490	12490	12867	12867
Adjusted R^2	0.107	0.107	0.050	0.050	0.282	0.282	0.077	0.077

Table IA5: Flow Performance Sensitivity

This table analyzes the sensitivity of fund flows to performance. For each year-month, fund alpha is calculated as the difference between realized excess return and predicted excess return based on the CAPM model, where the market factor loading is estimated using the previous 60 months of observations. Monthly alpha for each period is then computed accordingly. The independent variable, Alpha at time t, is calculated as a weighted average of returns from t – 1 to t – 3 (t – 18) using an exponential decay function with a decay parameter λ = 0.2. $\mathrm{Alpha}_{jt} = \frac{\sum_{s=1}^{3(18)} e^{-\lambda(s-1)}\tilde{\alpha}_{t-s}}{\sum_{s=1}^{3(18)} e^{-\lambda(s-1)}}$. Panel A reports results for Moral Funds and Benchmark Funds. Panel B reports results for Impact Funds and Benchmark Funds. Panel C reports results for Financial Funds and Benchmark Funds. Control variables include lagged log total AUM, log age, lagged expense ratio, and a load dummy. Standard errors are double clustered by fund and year-month.

Panel A: Moral Funds vs Benchmark Funds

	(1)	(2)	(3)	(4)
	Flow Rate	Flow Rate	Flow Rate	Flow Rate
Alpha*Moral Fund	-0.197***	-0.193***	-0.324***	-0.324***
	(-3.882)	(-3.839)	(-2.712)	(-2.751)
Alpha	0.371***	0.388***	0.815***	0.823***
	(10.368)	(10.808)	(12.039)	(12.303)
Moral Fund	0.301**	0.301***	0.269**	0.267**
	(2.620)	(2.628)	(2.069)	(2.059)
Flow Rate at t-4	0.023***	0.023***		
	(3.914)	(3.963)		
Flow Rate at t-19			0.049***	0.051***
			(4.963)	(5.272)
Year Month FE	No	Yes	No	Yes
Double Clustering at Fund Year Month	Yes	Yes	Yes	Yes
Alpha Decay Month	3	3	18	18
Observations	124968	124968	100845	100845
Adjusted R2	0.005	0.007	0.006	0.007

Panel B: Impact Funds vs Benchmark Funds

	(1) Flow Rate	(2) Flow Rate	(3) Flow Rate	(4) Flow Rate
Alpha*Impact Fund	-0.003	0.003	-0.010	-0.003
	(-0.038)	(0.030)	(-0.041)	(-0.014)
Alpha	0.371***	0.389***	0.814***	0.823***
	(10.360)	(10.817)	(12.033)	(12.315)
Impact Fund	0.503**	0.513**	0.385	0.394
	(2.148)	(2.185)	(1.503)	(1.521)
Flow Rate at t-4	0.023***	0.023***		
	(3.915)	(3.967)		
Flow Rate at t-19			0.049***	0.051***
			(4.869)	(5.161)
Year Month FE	No	Yes	No	Yes
Double Clustering at Fund Year Month	Yes	Yes	Yes	Yes
Alpha Decay Month	3	3	18	18
Observations	121202	121202	97446	97446
Adjusted R2	0.005	0.007	0.005	0.007

Panel C: Financial Funds vs Benchmark Funds

	(1) Flow Rate	(2) Flow Rate	(3) Flow Rate	(4) Flow Rate
Alpha*Financial Fund	-0.068	-0.049	-0.059	-0.049
	(-0.728)	(-0.525)	(-0.376)	(-0.313)
Alpha	0.371***	0.391***	0.814***	0.831***
	(10.323)	(10.985)	(11.974)	(12.391)
Financial Fund	0.167	0.186	0.136	0.150
	(1.358)	(1.578)	(1.119)	(1.243)
Flow Rate at t-4	0.022***	0.022***		
	(4.120)	(4.170)		
Flow Rate at t-19			0.050***	0.052***
			(5.544)	(5.821)
Year Month FE	No	Yes	No	Yes
Double Clustering at Fund Year Month	Yes	Yes	Yes	Yes
Alpha Decay Month	3	3	18	18
Observations	139755	139755	114384	114384
Adjusted R2	0.005	0.007	0.005	0.006

Table IA6: Moral Fund Votes on Environmental and Social Shareholder Proposals

This table reports the analysis of moral funds' votes on environmental and social shareholder proposals. For each fund-year, we calculate the percentages of Environmental and Social (ES) proposals received by the companies held by the fund for which the fund voted "For," "Against," "Abstain," and "Do Not Vote". Mean values are provided for pure moral, financial-moral, moral-impact, and financial-moral-impact funds.

	(1) Pure Moral	(2) Financial-Moral	(3) Moral-Impact	(4) Financial-Moral-Impact
Vote "For" ES	40.5%	54.2%	62.7%	79.2%
Vote "Against" ES	53.9%	41.4%	29.2%	10.5%
Abstain ES Vote	2.7%	1.7%	1.9%	2.5%
Do not vote	2.0%	2.0%	2.6%	0.1%
Observations	761	250	108	45

Table IA7: Difference-in-Differences Analysis of Emission Intensity

This table reports the dynamic treatment effect of being held by sustainable funds. We estimate a dynamic treatment effect model.

$$y_{i,t} = \sum_{k=-18}^{18} Holding_{i,t}^{k} + FEs + \varepsilon_{i,t}$$

Where $Holding_{i,t}^{k}$ is a dummy variable equal to one for firm i at calendar time t if t is k months before/after the first month in which firm i becomes held by a fund of a given sustainability type for $k \in [-18,18]$. For each treated firm—defined as a stock newly purchased by a sustainable fund—we identify a matched control firm based on ROA and log(total assets) within the same 2-digit NAICS sector. Control firms are further required to have received no sustainable fund investment between *t–24* and *t+24*. We define a fund's holding session as the uninterrupted period during which a stock appears in the fund's portfolio and the fund remains classified as a given type (impact, financial, or moral) at the time of entry. The session ends when either the stock disappears from holdings or the fund changes classification. Investment sessions lasting six months or fewer are excluded. Similar definitions apply across all fund types. The outcome variable is Scope 1&2 and Scope 1&2&3 carbon intensity. Firm and year-month fixed effects are included, and standard errors are clustered at the firm level. Columns (1) and (2), (3) and (4), (5) and (6), and (7) and (8) report results for impact, environmental impact, financial, and moral fund holdings, respectively. Environmental impact funds are defined as impact funds in the top quintile based on their support for environmental related shareholder proposals.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Impact		Env. Impact		Financial		Moral	
	1&2	1&2&3	1&2	1&2&3	1&2	1&2&3	1&2	1&2&3
t-18	2.815	1.954	2.678	2.485	0.873	0.791	2.509**	2.786**
	(0.914)	(0.607)	(0.648)	(0.598)	(1.069)	(0.923)	(2.101)	(2.213)
t-17	2.938	1.963	3.081	2.554	0.431	0.357	2.149*	2.716**
	(0.979)	(0.628)	(0.772)	(0.634)	(0.567)	(0.446)	(1.813)	(2.144)
t-16	2.836	1.934	2.378	2.166	-0.073	-0.164	2.179*	2.630**
	(0.972)	(0.640)	(0.614)	(0.549)	(-0.103)	(-0.218)	(1.877)	(2.123)
t-15	1.970	1.104	1.825	1.765	-0.503	-0.573	1.449	2.079*
	(0.679)	(0.369)	(0.492)	(0.464)	(-0.739)	(-0.775)	(1.262)	(1.679)
t-14	1.571	0.792	1.031	1.129	-0.727	-0.715	1.316	2.016
	(0.546)	(0.267)	(0.283)	(0.300)	(-1.081)	(-0.981)	(1.075)	(1.531)
t-13	-0.155	-0.713	0.757	0.993	-1.030	-1.081	1.081	1.666
	(-0.059)	(-0.259)	(0.209)	(0.264)	(-1.640)	(-1.554)	(0.895)	(1.285)
t-12	-0.372	-0.805	0.438	0.934	-1.315**	-1.317*	0.695	1.022
	(-0.139)	(-0.289)	(0.122)	(0.251)	(-2.089)	(-1.886)	(0.562)	(0.779)
t-11	-0.697	-1.311	-0.512	-0.064	-1.617**	-1.611**	0.649	0.962
	(-0.273)	(-0.489)	(-0.152)	(-0.018)	(-2.508)	(-2.265)	(0.561)	(0.781)
t-10	-1.023	-1.370	-1.262	-0.401	-1.780***	-1.789**	0.686	0.800
	(-0.398)	(-0.507)	(-0.381)	(-0.114)	(-2.705)	(-2.486)	(0.580)	(0.642)

t-9	-1.700	-2.084	-2.802	-1.846	-1.950***	-2.035***	0.145	0.409
	(-0.657)	(-0.768)	(-0.806)	(-0.498)	(-2.925)	(-2.835)	(0.121)	(0.325)
t-8	-1.762	-2.325	-2.880	-1.785	-2.235***	-2.193***	-0.217	-0.112
	(-0.707)	(-0.884)	(-0.846)	(-0.487)	(-3.295)	(-3.015)	(-0.177)	(-0.087)
t-7	-2.115	-2.726	-3.166	-1.878	-2.398***	-2.402***	-0.677	-0.442
	(-0.831)	(-1.013)	(-0.938)	(-0.513)	(-3.514)	(-3.301)	(-0.560)	(-0.350)
t-6	-2.696	-3.410	-3.834	-2.544	-2.449***	-2.523***	-1.107	-0.962
	(-1.070)	(-1.275)	(-1.133)	(-0.688)	(-3.533)	(-3.410)	(-0.897)	(-0.748)
t-5	-3.510	-4.212	-4.292	-3.084	-2.492***	-2.611***	-0.925	-0.814
	(-1.306)	(-1.486)	(-1.176)	(-0.779)	(-3.653)	(-3.565)	(-0.748)	(-0.631)
t-4	-3.730	-4.276	-5.054	-3.628	-2.794***	-2.780***	-1.204	-1.220
	(-1.394)	(-1.520)	(-1.381)	(-0.917)	(-4.034)	(-3.758)	(-0.980)	(-0.953)
t-3	-3.966	-4.485	-5.427	-4.187	-2.888***	-2.835***	-1.388	-1.369
	(-1.465)	(-1.574)	(-1.458)	(-1.039)	(-4.105)	(-3.766)	(-1.121)	(-1.063)
t-2	-4.007	-4.522	-5.559	-4.417	-2.715***	-2.564***	-1.276	-1.472
	(-1.476)	(-1.587)	(-1.508)	(-1.105)	(-3.806)	(-3.327)	(-0.993)	(-1.110)
t-1	-4.147	-4.584	-5.795	-4.510	-2.888***	-2.749***	-1.328	-1.418
	(-1.522)	(-1.593)	(-1.571)	(-1.122)	(-4.045)	(-3.563)	(-1.027)	(-1.055)
t	-3.349	-3.459	-5.540	-4.354	-3.028***	-2.785***	-1.682	-1.744
	(-1.237)	(-1.210)	(-1.453)	(-1.060)	(-4.122)	(-3.475)	(-1.299)	(-1.296)
t+1	-3.878	-4.067	-5.905	-4.896	-2.987***	-2.751***	-1.560	-1.500
	(-1.430)	(-1.426)	(-1.447)	(-1.125)	(-3.963)	(-3.366)	(-1.190)	(-1.102)
t+2	-3.668	-3.962	-5.833	-4.690	-2.935***	-2.750***	-1.627	-1.497
	(-1.368)	(-1.397)	(-1.425)	(-1.081)	(-3.849)	(-3.368)	(-1.244)	(-1.095)
t+3	-3.354	-3.838	-5.475	-4.627	-2.907***	-2.798***	-1.810	-1.680
	(-1.254)	(-1.359)	(-1.312)	(-1.050)	(-3.708)	(-3.323)	(-1.367)	(-1.209)
t+4	-3.424	-3.864	-6.113	-5.029	-2.790***	-2.616***	-1.609	-1.500
	(-1.275)	(-1.356)	(-1.453)	(-1.135)	(-3.403)	(-2.968)	(-1.235)	(-1.098)
t+5	-3.487	-3.749	-6.763	-5.440	-2.804***	-2.651***	-1.661	-1.574
	(-1.286)	(-1.310)	(-1.579)	(-1.213)	(-3.373)	(-2.971)	(-1.250)	(-1.138)
t+6	-3.995	-4.223	-7.183	-5.902	-2.499***	-2.337***	-1.808	-1.598
	(-1.461)	(-1.471)	(-1.634)	(-1.286)	(-2.966)	(-2.593)	(-1.367)	(-1.162)
t+7	-4.288	-4.559	-8.129*	-6.843	-2.093**	-1.798*	-1.427	-0.994
	(-1.449)	(-1.487)	(-1.851)	(-1.508)	(-2.227)	(-1.772)	(-1.016)	(-0.686)
t+8	-4.367	-4.657	-7.704*	-6.391	-1.996**	-1.563	-1.269	-0.768
	(-1.461)	(-1.510)	(-1.751)	(-1.399)	(-2.117)	(-1.539)	(-0.881)	(-0.516)
t+9	-4.737	-4.843	-8.821*	-7.273	-1.602	-1.177	-1.340	-0.945
	(-1.539)	(-1.534)	(-1.931)	(-1.550)	(-1.639)	(-1.119)	(-0.909)	(-0.623)
t+10	-5.338*	-5.829*	-8.714*	-7.274	-1.275	-0.719	-1.242	-0.920
	(-1.738)	(-1.822)	(-1.908)	(-1.531)	(-1.308)	(-0.679)	(-0.838)	(-0.604)
t+11	-5.259	-5.778*	-9.426**	-7.604	-1.122	-0.646	-1.002	-0.718
	(-1.632)	(-1.724)	(-1.998)	(-1.554)	(-1.111)	(-0.590)	(-0.658)	(-0.460)
t+12	-5.654*	-6.148*	-9.953**	-8.299*	-1.130	-0.669	-1.139	-0.982
	(-1.702)	(-1.781)	(-2.055)	(-1.662)	(-1.064)	(-0.583)	(-0.712)	(-0.599)

t+13	-7.187**	-7.470**	-10.544**	-9.118*	-0.825	-0.170	-0.698	-0.034
	(-2.084)	(-2.100)	(-2.152)	(-1.802)	(-0.778)	(-0.149)	(-0.434)	(-0.020)
t+14	-7.096**	-7.289*	-11.145**	-9.044*	-0.685	-0.132	-0.760	0.072
	(-1.972)	(-1.957)	(-2.326)	(-1.840)	(-0.616)	(-0.109)	(-0.449)	(0.041)
t+15	-6.428*	-6.392*	-10.756**	-8.455*	-0.483	0.102	-0.878	-0.117
	(-1.725)	(-1.656)	(-2.146)	(-1.676)	(-0.414)	(0.081)	(-0.501)	(-0.065)
t+16	-6.614	-6.998*	-11.989**	-9.637*	-1.833*	-1.896	-1.500	-0.370
	(-1.615)	(-1.657)	(-2.389)	(-1.863)	(-1.678)	(-1.597)	(-0.875)	(-0.210)
t+17	-6.554	-6.989	-11.544**	-8.912*	-1.591	-1.488	-1.207	-0.160
	(-1.509)	(-1.567)	(-2.229)	(-1.681)	(-1.366)	(-1.173)	(-0.666)	(-0.086)
t+18	-6.283	-6.595	-12.036**	-8.945	-0.903	-0.708	-1.078	0.359
	(-1.380)	(-1.409)	(-2.183)	(-1.596)	(-0.717)	(-0.518)	(-0.610)	(0.196)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year-Month FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Cluster at Firm	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	111661	111661	65717	65717	230796	230796	254848	254848
Adjusted R2	0.953	0.959	0.956	0.964	0.930	0.936	0.930	0.937

Figure IA1: Growth of Sustainable Mutual Funds, 2014–2023

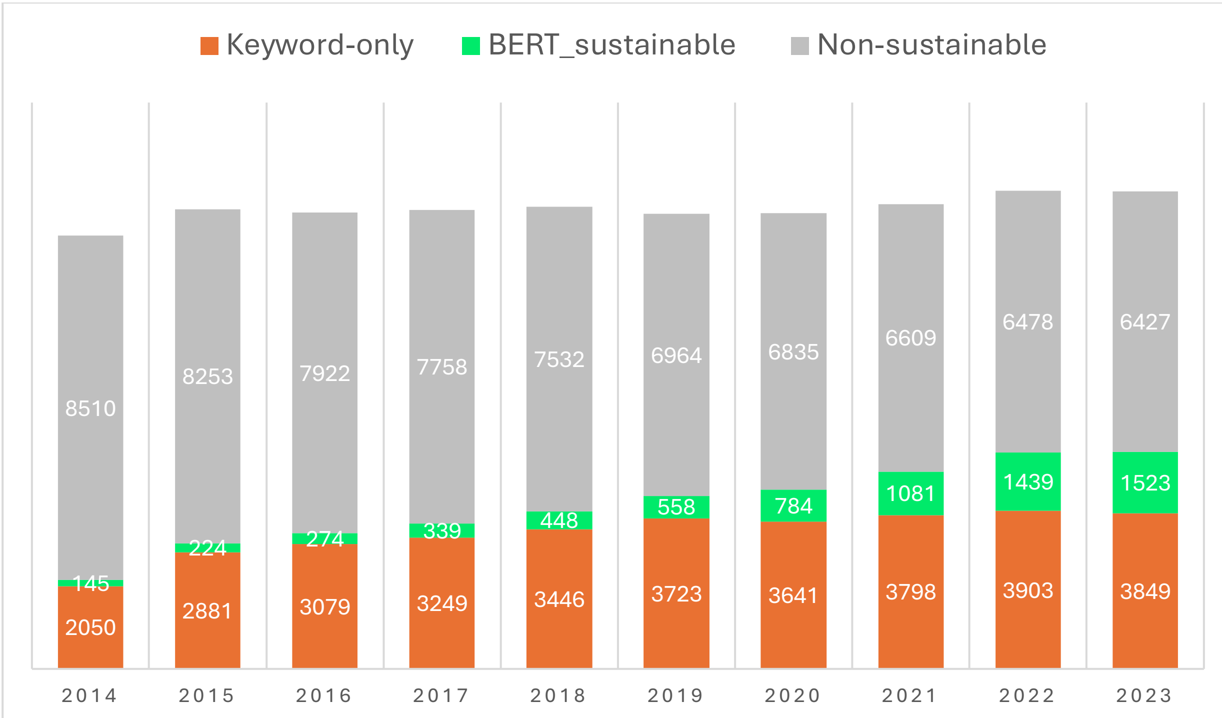

Panel A: Fund Count

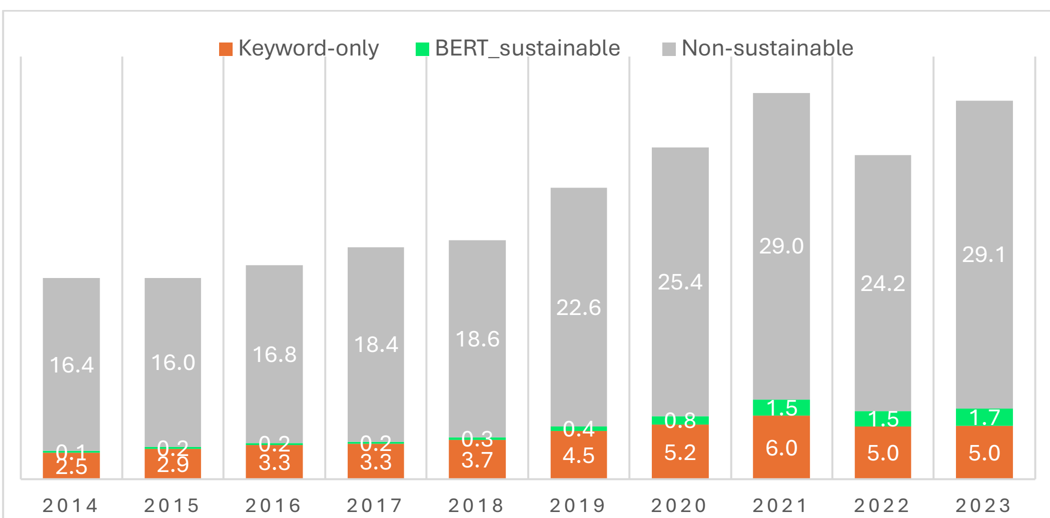

Panel B: Fund AUM ($T)

This figure plots the number of funds and total assets under management (AUM) from 2014Q2 to 2023Q4 for three groups: sustainable funds, keyword-only funds, and non-sustainable funds. Sustainable Funds are funds that the BERT model identifies as financial, moral, or impact. "Keywords-only" refers to funds that mention at least one keywords related to sustainability in the "Principal Investment Strategy" section of its prospectus, but that the BERT model identifies zero sentences as meeting the sustainability criteria. "Non-sustainable" refers to funds that do not mention any sustainability-related keywords.

Figure IA2: Distribution of Sustainable Fund Classifications

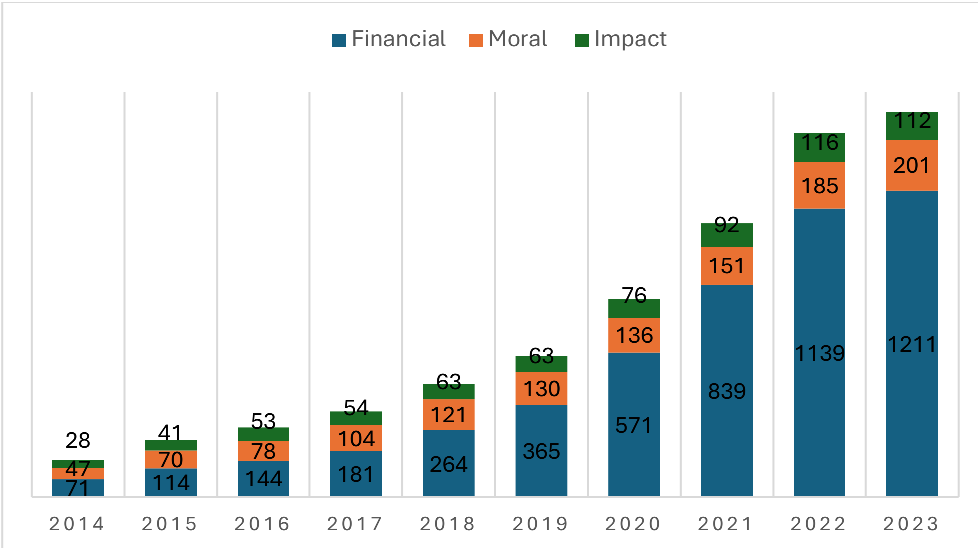

Panel A: Fund Count

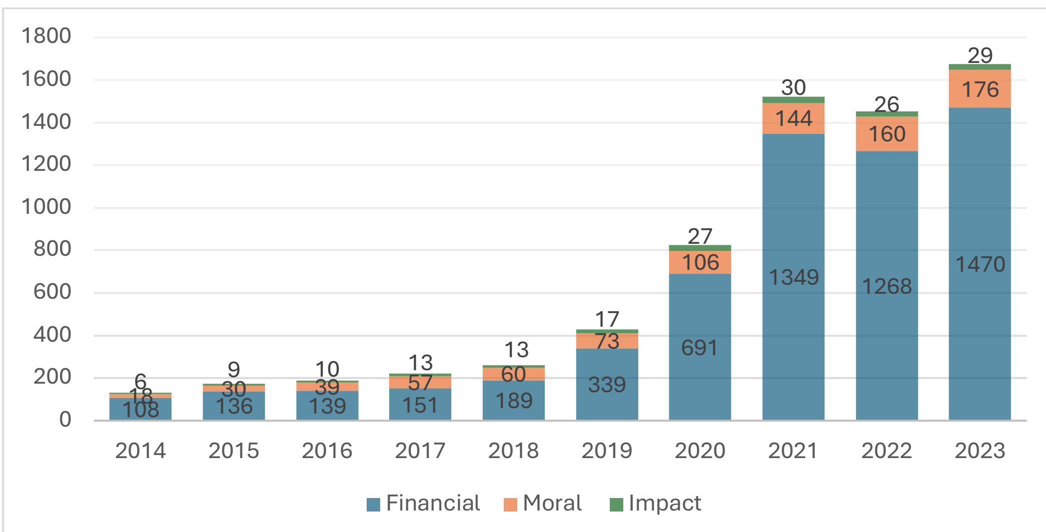

Panel B: Fund AUM ($B)

This figure plots the cross-sectional distribution of sustainable fund classifications by year. Classifications are based on the % of sustainable sentences for each fund that are classified as financial, moral, and impact. A fund is assigned the category with the largest % share of all sustainable sentences. When there is a two-way (three-way) tie for the largest % share, each category receives 50% (33%) of the share count/AUM of the fund. The fund count sums up to 146 in 2014 due to rounding, as 2 funds are dual tie-funds and are assigned 0.5 fund count each in 2 categories.

Figure IA3: Decompose AUM

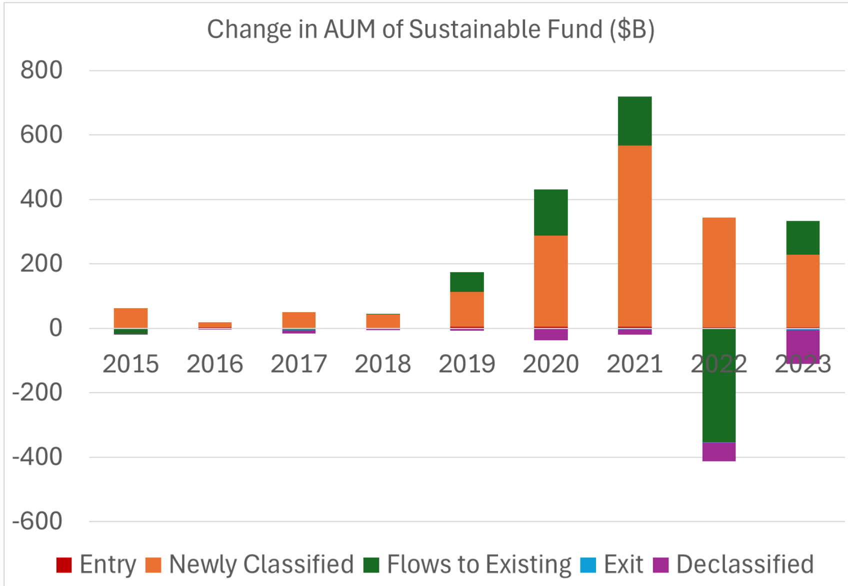

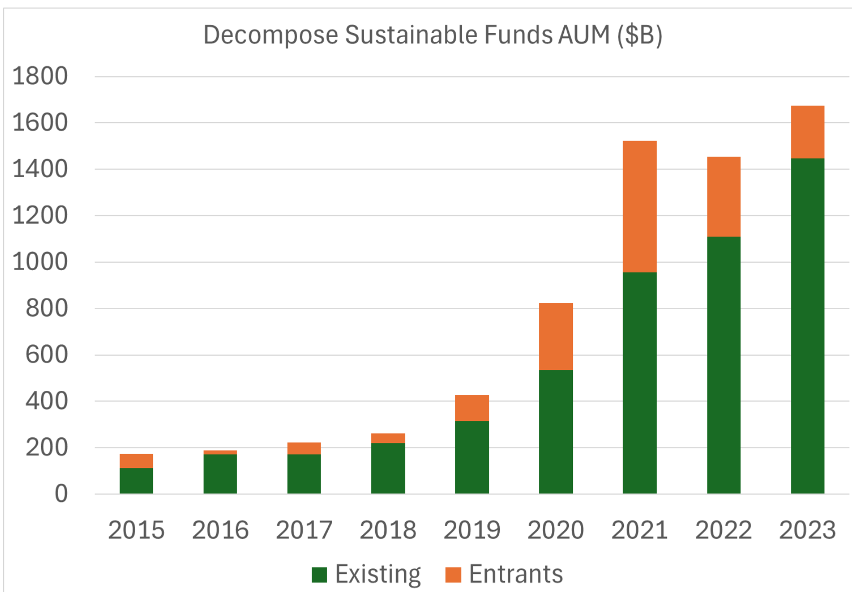

This figure plots the annual change and decomposition of assets under management (AUM) for sustainable funds from 2015 to 2023. In Panel A, for each year we compute the change in AUM from the previous year to the current year and decompose this change into several components: increases in AUM due to newly created funds classified as financial (Entry), due to funds that are newly classified (Newly Classified), changes in AUM for previously classified funds (Flows to Existing), decreases in AUM due to fund exits (Exit), and decreases in AUM due to funds being declassified (Declassified). In Panel B, we decompose total AUM in a given year into AUM from funds that were classified as financial in previous years (Existing) and AUM from funds that either entered in that year or were newly classified (Entrants).

Figure IA4: Robust DiD Analysis of Firm Emission Intensity during Holding Periods

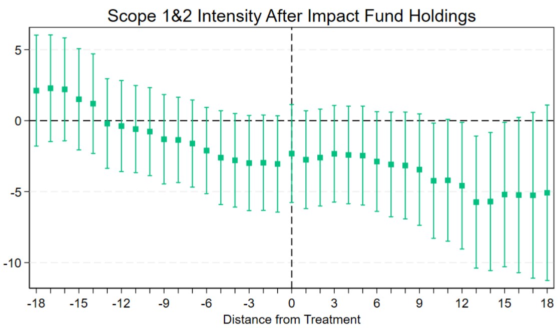

Panel A: Impact Fund Holdings

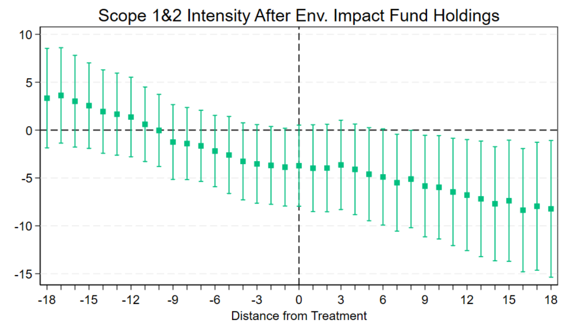

Panel B: Environmental Impact Fund Holdings

This figure shows the dynamic treatment effect on firms' Scope 1 & 2 carbon intensity after acquisition by sustainable funds. We estimate a dynamic treatment effect model.

$$y_{i,t} = \sum_{k=-18}^{18} Holding_{i,t}^{k} + FEs + \varepsilon_{i,t}$$

where $Holding_{i,t}^{k}$ is a the total number of simultaneous holding events for firm i at calendar time t if t is k months before/after the first month in which firm i becomes held by a fund of a given sustainability type for $k \in [-18,18]$. For each treated firm—defined as a stock newly purchased by a sustainable fund—we identify a matched control firm based on ROA and log(total assets) within the same 2-digit NAICS sector. Control firms are further required to have received no sustainable fund investment between t–24 and t+24. The outcome variable is Scope 1&2 carbon intensity. Firm and year-month fixed effects are included, and standard errors are clustered at the firm level. Confidence intervals are reported at the 90% level. Panel A reports results for impact fund holdings, Panel B reports environmental impact fund holdings, focusing on funds in the top quintile based on their support for environmental related shareholder proposals.